Next Generation Internet

Implementation Plan
Draft

July 1997

Large Scale Networking
Next Generation Internet Implementation Team

Note: This plan incorporates the comments received from the Presidential Advisory Committee on High Performance Computing and Communications, Information Technology, and the Next Generation Internet; Members of Congress and their staff; an NGI workshop sponsored by Computer Research Association, Computer Systems Policy Project, and Cross Industry Working Team; industry; academia; and the public. Please note that both this document and the NGI Concept Paper are based upon the Presidential requested level of funding. Congressional action may result in changes that will be incorporated into these documents after final FY98 budget approval.

Comments are always encouraged. Please send them to ngi@ccic.gov or fax them to 703-306-4727. If you need additional information, please contact the National Coordination Office at 703-306-4722.

Acknowledgments

The following people contributed to this plan:

Michael Ackerman, NIH
Robert Aiken, DoE
Debra Bailey, NASA
Richard desJardins, NASA
Richard DuBois, NIH
Phil Dykstra, DoD
Don Endicott, DoD
Christine Falsetti, NASA
Jim Fowler, NIST
Ken Freeman, NASA
Bert Hui, DARPA
Gary Koob, DARPA
Mark Luker, NSF
Doug Montgomery, NIST
Hilarie Orman, DARPA
Alex Poliakoff, Dept. of Education
Mary Anne Scott, DoE
George Seweryniak, DoE
Carl Stanton, NOAA
Dave Staudt, NSF
Bill Turnbull, NOAA

CONTENTS

List of Figures

1. EXECUTIVE SUMMARY

The U.S. Government's Internet investments have been incredibly successful. The Internet has grown at nearly 100 percent per year since 1988 and Internet traffic has been growing at a rate of 400 percent per year recently. The Internet has created jobs and whole new industries. American business and government organizations are increasingly dependent on the Internet.

There are signs of frailty and limitations in the Internet's current capabilities. This frailty is being exacerbated by high bandwidth access and multimedia applications. The solutions to these problems are beyond the scope of any single institution, company or industry. The Next Generation Internet (NGI), with its broad agenda and ability to involve government, research institutions, and the business sector, is a timely program that will address these problems.

By the year 2000, more than 50 percent of the U.S. population will have access to the Internet. If the coming exponential improvements in computing and communications are to benefit American industry and the public we must make a few key strategic investments in research and development now. NGI is one of those key strategic investments. NGI will help to create the environment in which R&D breakthroughs are possible.

The federal government has a unique role to play in stimulating progress. The NGI initiative, together with other investment sectors shown in figure 1, will create the foundation for the networks of the 21st century, setting the stage for networks that are much more powerful and versatile than the current Internet. The NGI will foster partnerships among academia, industry, and government that will keep the United States at the cutting-edge of the information and communications technologies. The NGI will also stimulate the introduction of new multimedia services in our homes, schools, and businesses as the technologies and architectures designed and developed as part of the NGI are incorporated into products and services that are subsequently made available to the general public. The NGI program is essential to sustain U.S. technological leadership in computing and communications and enhance U.S. economic competitiveness and commercial leadership.

This initiative is possible only because of the very strong agency programs that are currently under way. The Large Scale Networking R&D crosscut for FY 1998, for example, is $288.3 million, which includes $100 million for NGI. The present document focuses directly on the goals and objectives of the $100 million NGI initiative.

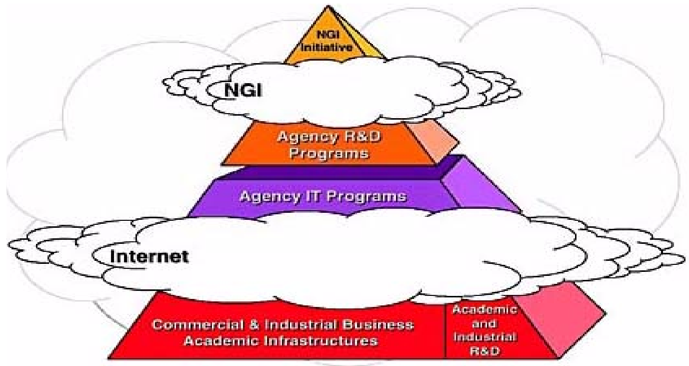

Figure 1. Next Generation Internet programs.

Specifically, the NGI has three goals as outlined in the NGI Concept Paper (see *http:// www.ngi.gov/*). They are as follows:

1. Promote experimentation with the next generation of networking technologies.
2. Develop a next generation network testbed to connect universities and Federal research institutions at rates that demonstrate new networking technologies and that support future research.
3. Demonstrate new applications that meet important national goals and missions and that rely on the advances made in goals 1 and 2.

The Administration has made an initial funding commitment of $100 million per year for 3 years or a total of $300 million. To achieve the above goals, the NGI will be built on the base of current R&D activities and programs in the participating Federal agencies. Furthermore, it will call on substantial matching funds from its private sector partners and will collaborate with academia.

This document supplements the NGI Concept Paper and provides additional details on methods and steps proposed to reach those goals.

1.1 GOAL 1: EXPERIMENTAL RESEARCH FOR ADVANCED NETWORK TECHNOLOGIES

Goal 1 activities will focus on research, development, deployment, and demonstration of the technologies necessary to permit the effective, robust, and secure management, provisioning, and end-to-end delivery of differentiated service classes. These activities cluster into three major tasks: (1) network growth engineering,(2) end-to-end quality of service (QoS), and (3) security.

The high speed and advanced communications capacity (developed under Goals 2.1 and 2.2) will enable advanced applications for the Department of Defense (DoD), the Department of Energy (DoE), the National Aeronautics and Space Administration (NASA), the National Science Foundation (NSF), and other agency users. However, increased bandwidth alone will not be sufficient to meet the dependability, classes of services, security, and real time demands of emerging and next generation applications, such as collaboration, wide area distributed computing, and teleoperation and control.

The challenge for Goal 1, then, is to develop the advanced capabilities required for the Goal 2 testbed networks and ensure that the advanced capabilities of Goal 2 networks can be made predictably and reliably accessible to a broad range of users sharing a common infrastructure. This will involve Goal 1 technologies being aggressively deployed into the Goal 2 networks. Since the Goal 2 testbeds may incur temporary service degradations as a result of the experimental alpha deployment of Goal 1 technologies and Goal 3ís use of these technologies, they must be considered in applications planning.

This will be joint agency effort with the Defense Advanced Research Projects Agency (DARPA) as the lead and with the participation of DoE, the National Institute of Standards and Technology (NIST), NASA, NSF, and other agencies.

1.2 GOAL 2: NGI FABRIC

The networks developed under the NGI initiative will connect at least 100 sites—universities, Federal research institutions, and other research partners—at speeds 100 times faster than those of today’s Internet, and will connect on the order of 10 sites at speeds 1,000 times faster than the current Internet.

This goal addresses end-to-end connectivity (to the workstation) at speeds from 100+ million bits per second (Mbps) up to 1+ billion bits per second (Gbps.) Some networks have already achieved OC-12 (Optical Carrier) speeds (622 Mbps) on their backbone links and some experimental links are running at 1+ Gbps, but end-to-end usable connectivity is typically limited to less than 10 Mbps because of bottlenecks or incompatibilities in switches, routers, local area networks, and workstations. Goal 2 addresses these shortcomings by developments and demonstrations involving two subgoals. Goal 2 testbed network fabric will aggressively incorporate the technologies developed in goal 1.

1.2.1 Subgoal 2.1: High Performance Connectivity

The Goal 2.1 demonstration network fabric will function as a distributed laboratory. It will deliver a minimum improvement of 100 times (or greater) over the current Internet performance on an end-to-end basis to at least 100 interconnected NGI participating universities, national laboratories, and Federal research sites demonstrating research and other important applications that require such an infrastructure. This network fabric will be large enough to provide a full system, proof-of-concept testbed for hardware, software, protocols, security, and network management that is required in the commercial NGI. This goal will address not only accessible, but also remote sites and EPSCoR (Experimental Program to Stimulate Competitive Research) states. Experiments are anticipated to assist research in reaching beyond the current Internet infrastructure.

Goal 2.1 is a joint agency effort led by DoE, NSF, and NASA, with participation from DoD and other agencies.

1.2.2 Subgoal 2.2: NGI Technologies and Ultra High Performance Connectivity

Goal 2.2 addresses the development of ultrahigh speed switching and transmission technologies and of end-to-end network connectivity at 1+ Gbps. Because of its high risk and pioneering nature, networks involved will be limited initially to approximately 10 NGI sites, and only a limited number of applications will be implemented. Some of the nodes of Goal 2.2 will overlap with those of Goal 2.1.

Attainment of this goal, together with the technologies developed in Goal 1, will lay the groundwork for terabit per second (Tbps) networks operated by appropriate network management and control and guaranteeing end-to-end quality of service. Working in partnership with industry is the key to a shared infrastructure that can be used profitably to support high end scientific users and large numbers of commercial users.

Goal 2.2 is a joint agency effort with the DARPA as the lead, and with the participation of DoE, NASA, NSF, and other government agencies.

1.3 GOAL 3: REVOLUTIONARY APPLICATIONS

To achieve Goal 3, the participating Federal agencies established procedures to identify appropriate applications to be tested. These applications require the advanced networking capabilities of Goals 1 and 2, and agencies must be willing to adapt their applications to take advantage of these capabilities. The resulting NGI applications will integrate advanced networking and application technologies.

A coordinated selection process will be used to ensure that applications tested and demonstrated on the NGI network provide robust, realistic, complete tests of technologies that are extensible and adaptable to other applications. The selection of NGI applications is an iterative process with Federal, academia, and industry participation. Applications will be derived from the Federally focused applications in appropriate technology classes, for example, digital libraries, remote operation of medicine, environment, crisis management, manufacturing, basic sciences, and Federal information services.

This joint agency effort will be coordinated among the participating agencies. Since most of the funding for applications will come from the applications themselves, leadership will be provided by means of domain specific affinity groups. Participation will be encouraged from a broad range of agencies with demanding networking applications. Applications will also be solicited from other interested research entities within academia and industry.

1.4 RESOURCES

The specific skills and experience that the participating agencies bring to the initiative provide the essential base upon which it is built. Because of this strong base, the initiative is expected to succeed; without that base the initiative would involve much more risk. Specific agency strengths include:

1. DARPA: Pioneer of long term networking research; developed cutting-edge network technologies; great strength in network management and services.
2. DoE: Long term experience in managing production and research networks; specialized skills in networking technology; great strength in mission driven applications and in system integration.
3. NASA: Experience in network management and in specialized network testbeds; strength in mission driven applications involving high data rates; great strength in system engineering and integration.
4. NSF: Special relationships with the academic community; experience in network research and in managing networks; great strength in scientific applications.

5. NIST: long experience in standards development, networking research, metrology, computer systems security, systems integration for manufacturing applications, and in testbeds involving many industrial partners.

6. National Library of Medicine (NLM)/National Institute of Health (NIH): Extensive experience in medical research; great strength in health care applications.

NGI FY 1998 Proposed $105* Million Budget, $ million

Goal	DoD/DARPA	NSF	DoE	NASA	NIST	NLM/NIH	Total
Goal 1: Experimental Research	20	2	6	2	3		**33**
Goal 2: Next Generation Network Fabric	20	7	25	3			**55**
Goal 3: Revolutionary Applications		1	4	5	2	5	**17**
Total	**40**	**10**	**35**	**10**	**5**	**5**	**105**

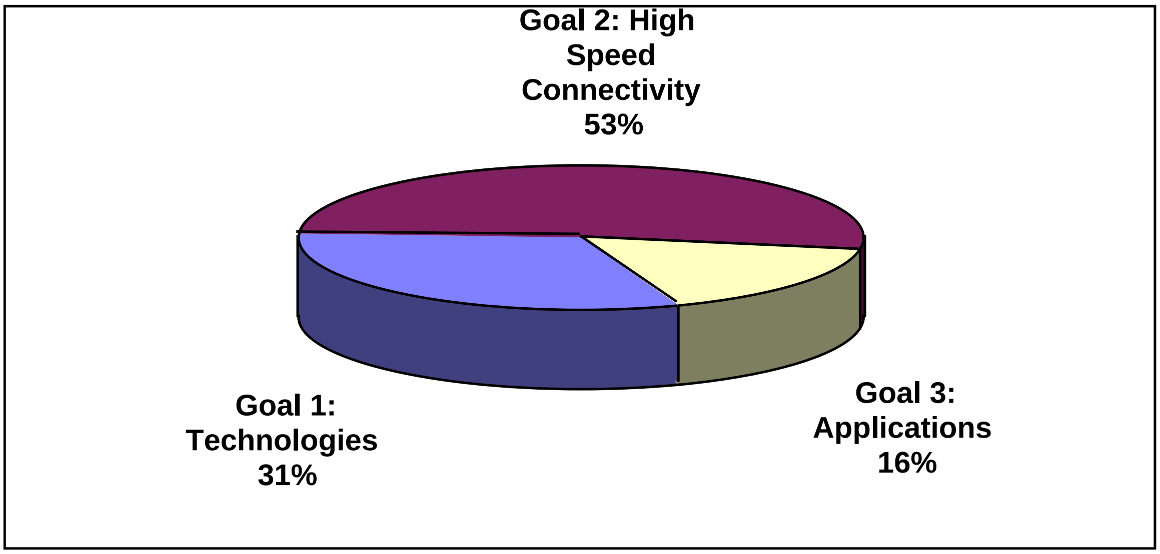

Figure 2. NGI FY 1998 Funding by Goal.

(Note that the initiative was originally proposed at $100 million per year with funding expected from additional agencies who want to be part of the program. NLM/NIH is the first example of an additional NGI initiative partner. Therefore, the total proposed budget is now $105 million allocated as shown in figure 2.)

1.5 MANAGEMENT

An NGI implementation team will be established to coordinate research agendas across all goals. The NGI will be under the overall guidance and direction of the NGI Implementation Team (NGI IT) reporting to the Large Scale Networking (LSN) working group as detailed in the organization chart shown in figure 3.

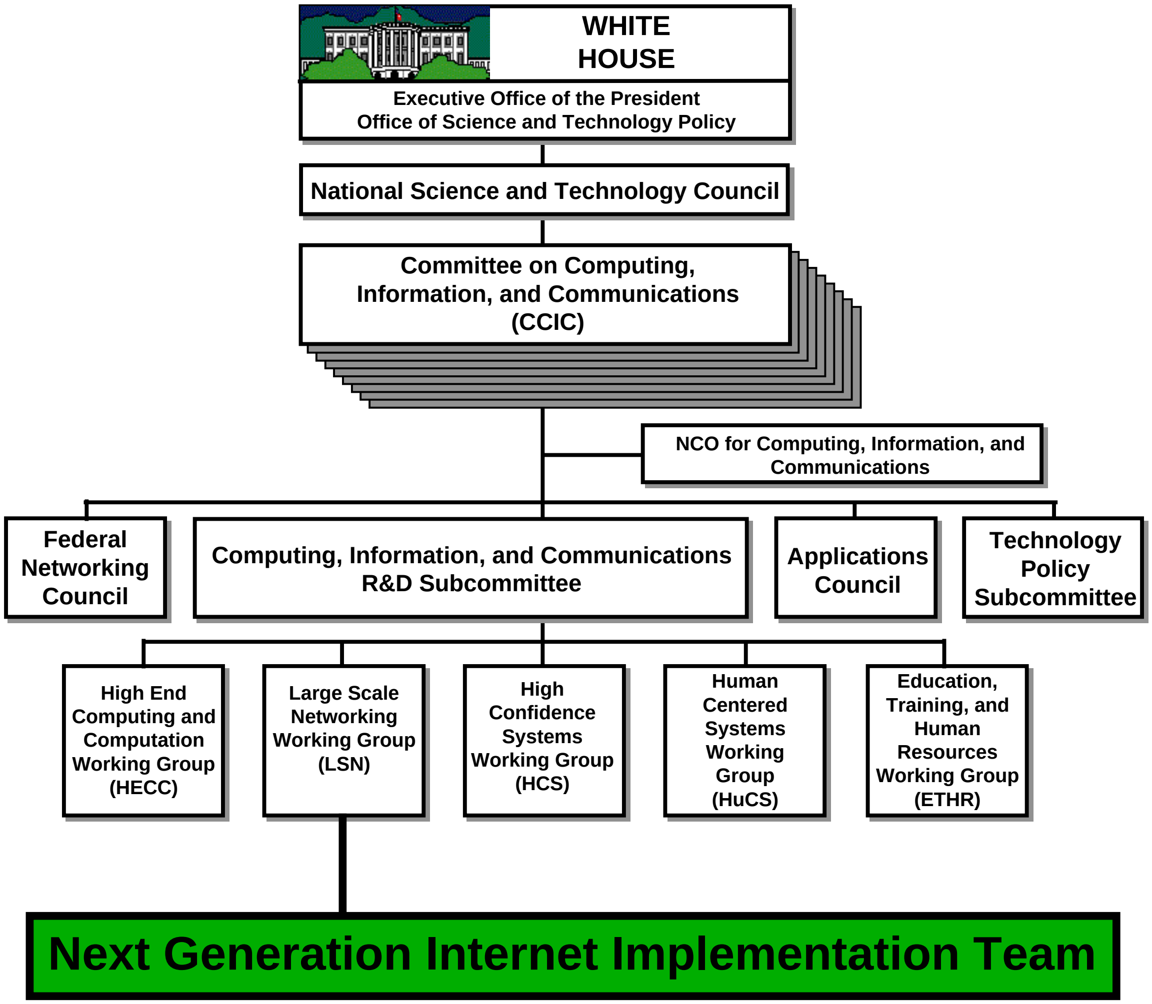

Figure 3. NGI Implementation Team reporting structure.

1.5.1 Coordination

The team will include appropriate agency program managers, and will include experts from academia, industry, and Federal laboratories.

The Implementation Team will meet as often as necessary to accomplish these goals, but at least four times per year after the initial implementation has been accepted. The detailed implementation plans for each of the goals are described in this plan.

1.5.2 NGI Participation

Each agency will use its own method for soliciting calls for research (e.g., solicitations, broad area announcements, calls for proposals) for all goals, and will ensure coordination among the agencies. This will be accomplished through the use of other agency program managers and experts as reviewers of resulting proposals, as well as through interagency program manager coordination activities such as those of the NGI Implementation Team.

Calls for proposals will be issued at least once at the beginning of each fiscal year. The primary selection criteria for NGI sites for Goal 2.1 will be based primarily on a siteís ability to demonstrate an NGI class application and the site’s use of the technologies in Goals 1, 2.1, and 2.2. The sites will also be required to demonstrate that they possess the expertise and infrastructure necessary to demonstrate these applications on an end-to-end basis. Sites that do not possess NGI applications, but do possess the necessary technology (i.e., Goals 1 and 2) and expertise may also be considered for the proposed award.

1.6 TIME LINE SUMMARY

Initiative Deliverables	First Achieved
1. 100+ site high performance testbed providing OC-3 (155 Mbps) connections over OC-12 (644 Mbps) infrastructure	1999
2. Federal, academic and industry partnerships begin conducting applications/networking research on the 100x testbed	1999
3. 10+ site ultrahigh performance test bed providing OC-48 connections (2.5 Gbps)	12000
4. Consortium conducting networking/applications research on the 1000x testbed	2001
5. Tested models for NGI protocols, management tools, QoS provisions, security, and advanced services	2000
6. 100+ high value applications testing and benefiting from high performance testbed (e.g., remote, real time,	2000

collaborative NGI network control of selected laboratories)

7. Integrate the two split network configurations to validate it reliable data communications 2001

8. Terabit-per-second packet switching demonstrated 2002

9. 10+ advanced applications testing and benefiting from ultra high performance testbed 2002

Note that dates presented in the remainder of the document sometimes span multiple years because the milestone or deliverable may evolve over time from base implementation to usable tools or capabilities. In each case, the beginning date identifies when the first instance of a tool or capability is expected, and the last date indicates when a fairly robust usable version is expected. Where appropriate, yearly demonstrations may also be conducted. Dates appearing in this document may also fluctuate as a result of the research proposals received in response to solicitations and as a result of similar work at other agencies.

2. GOAL 1: EXPERIMENTAL RESEARCH FOR ADVANCED NETWORK TECHNOLOGIES

2.1 INTRODUCTION AND STRATEGY

The Next Generation Internet (NGI) program will be the pathway to terabit-per-second network speeds over wide area advanced capability networks. This will be done by means of a partnership with industry leading to a shared infrastructure that can be used profitably to support high end users, as well as large numbers of typical commercial users. Although the high speed communication capability (developed under Goal 2.2) will enable advanced applications for the Department of Defense (DoD), the Department of Energy (DoE), the National Aeronautical and Space Administration (NASA), the National Science Foundation (NSF), and other agency users, increased bandwidth alone will be insufficient to meet the dependability, security, and real time demands of emerging and next generation applications, such as collaboration, wide area distributed computing, and teleoperation and control. The challenge for Goal 1, then, is to ensure that the advanced capabilities of both Goal 2.1 and Goal 2.2 networks can be made predictably and reliably accessible to a broad range of users sharing a common infrastructure.

Goal 1 activities will therefore focus on multiagency coordinated development, deployment, and demonstration of the technologies necessary to permit the effective, robust, and secure management, provisioning, and end-to-end delivery of differentiated service classes. These activities cluster into three major tasks: network growth engineering, end-to-end quality of service (QoS), and security.

This will be a joint agency effort with the Defense Advanced Research Projects Agency (DARPA) as the lead with augmenting and complementing participation by DoE, the National Institute of Standards and Technology (NIST), NASA, NSF and other agencies. Each agency will encourage participation in these research areas through its normal mechanisms (e.g., solicitations, broad agency announcements). The coordination of the resulting research proposals will be accomplished through cross agency participation in review panels and coordination by the agency program managers to ensure that proposals do not unnecessarily duplicate other efforts. The intent is to focus the research so that the total of the parts leads to an integrated solution.

The NGI strategy focuses on developing the most important and highly leveraged aspects of internetworking technology: network growth engineering, end-to-end QoS, and security.

2.1.1 Network Growth Engineering

The goals of this task are to (1) create and deploy tools and algorithms for planning and operations that guarantee predictable end-to-end performance at scales and complexities of 100 times those of the current Internet; (2) facilitate management of large scale internetworks operating at gigabit to terabit speeds supporting a range of traffic classes on a shared infrastructure; and (3) create an infrastructure partnership through which lead users (government and research) share facilities with the general public, thereby accelerating the development and penetration of novel network applications.

This task will develop and integrate technologies for network planning and simulation; network monitoring, analysis, and control; innovative data delivery; and shared infrastructure management for lead users. The highly automated services envisioned in this task lead to the goal of building strong security mechanisms into the components.

2.1.1.1 Planning and Simulation

Planning large network interconnections is now primarily a manual process that is not tied to any runtime tools or distributed efforts. Under this subtask, a network planning description language will be developed as a community standard facilitating not only the initial planning but also maintenance of requirements throughout the network life cycle.

Metrics: Ability to plan, coordinate, and maintain 100 organizational networks.

2.1.1.2 Monitoring, Control, Analysis, and Display

Network engineering and management requires tools for gathering data and information, analyzing it, and issuing control commands based on the results of the analysis. Current practice uses planned analysis based on protocol headers and aggregated statistics. This subtask will take on the challenge of presenting runtime analysis based on distributed communication patterns and very high communication speeds that would overwhelm current tools.

Metrics: Demonstrate monitoring and analysis at speeds of Optical Carrier (OC)-48 (2.5 Gbps) and higher.

2.1.1.3 Integration

The goal of this subtask is to ensure that the network engineering tools and the high performance transmission and switching technologies work smoothly together for accomplishing end-to-end management of leading edge user requirements. The requirements will be developed with the users, and the tools will be integrated into the NGI testbed, developing 10 distributed management stations with monitoring and command interfaces to all connected equipment; half of the management stations will be in the Goal 2.1 testbed and half will be in the Goal 2.2 testbed.

Metrics: Demonstrate 25 percent utilization improvement in Goal 2.1 network over 3 months and 100 percent improvement in Goal 2.2 network over 3 months.

2.1.1.4 Data Delivery

The management software will work in a tightly bound interlock with new strategies for managing and controlling data delivery in networks. This subtask will develop network interior nodes that combine methods previously seen as disjoint or mutually exclusive: routing and switching, best effort and priority traffic, dynamic routing and virtual circuits, greedy admission versus guaranteed delivery, and flat-rate versus variable costing. Tools that permit network engineers to adjust the strategy trade-offs to best meet their requirements will be prototyped and tested in the high speed arena.

Metrics: Demonstrate 100 percent improvement in throughput using heterogeneous, application specific routing strategies.

2.1.1.5 Managing Lead User Infrastructure

The DoD, DoE, NASA, NSF, and other government agencies, as well as the research community at large, typically have lead user requirements for telecommunication facilities that require speed and complexity that are orders of magnitude beyond those of the typical users. This task will investigate architectural concepts, management strategies, and operational arrangements that will allow lead users to concurrently share the same infrastructure as conventional users at a variety of levels. Extension of this dual modality into the campus infrastructure for end-to-end support will also be investigated.

Metrics: Demonstrate lead user support by striping, that is, partitioning a data stream across multiple low bandwidth channels to emulate a high bandwidth channel, over 100 ordinary channels without performance loss.

2.1.2 End-to-End Quality of Service

The goals of this task are to facilitate the delivery of end-to-end ensured QoS to applications and to ensure that these technologies can be tailored for use by and made available to lead users who have demanding requirements. The strategy is to allow users to negotiate application specific trade-offs among such parameters as bandwidth, latency, precision, and reliability in order to obtain predictable performance at a known quality level. Exploiting emerging network level mechanisms is difficult, however, as they are semantically far removed from the applications they are intended to support and are accessible only through many layers of software.

End-to-end QoS assurance requires an approach that spans these layers of operating system and middleware in order to effectively deliver network level QoS guarantees. This task will develop and demonstrate a comprehensive QoS management architecture, drill down technologies to facilitate propagation of QoS constraints across software and network layers, and next generation network technologies to support QoS. Research addressing issues specific to wireless networking and nomadicity are beyond the scope of this initiative; however, the QoS framework will be general enough to accommodate the integration of wireless services and the eventual integration of nomadicity support. This effort will be fully coordinated with on-going agency and industry advances in mobile networking and computing.

2.1.2.1 Baseline Quality of Service Architecture

The baseline QoS management architecture will provide the framework of models, languages, and protocols to permit distributed applications to specify multidimensional QoS requirements, negotiate acceptable trade-offs and confidence levels, and receive feedback on delivered QoS enabling adaptation. Application Program Interfaces (APIs) supporting the propagation of QoS constraints and feedback through software layers will be developed, as well as the requisite admission control, accounting/costing, security, and priority mechanisms both within the network infrastructure and the periphery ingress/egress environments.

Metrics: Demonstrate ability to handle differentiated service classes and to reduce variance in end-to-end performance by factors of three to five for multimedia traffic.

2.1.2.2 Drill Down Technologies

Current technologies support composition of *functionality* across system layers but not the composition of their QoS properties. This subtask will develop technologies that can be used to drill down and expose interfaces to QoS and network management capabilities that are presently hidden within the individual layers. Emphasis will be on techniques appropriate for operating systems, communications libraries, and middleware services, including distributed objects, and providing direct access from the applications to network layer components and objects.

Metrics: Demonstrate factor of 3 to 5 reduction in communications overhead attributable to systems software.

Next Generation Network Technologies

Next generation network technologies, particularly those supporting QoS, are critical to the success of the Internet, but they lack an adequate experience base at the speeds and scale envisioned for NGI. This subtask will accelerate deployment of and experimentation with technologies such as class based queuing, Resource Reservation Protocol (RSVP), and Internet Protocol Version 6 (IPv6), as well as develop new technologies for admission control, accounting/costing, scheduling, prioritization in both IP and Asynchronous Transfer Method (ATM) networks.

2.1.3 Security

Security's essential role in the NGI is to support several goals: a secure and fair means of user access to and use of network resources (e.g., QoS), smart network management, internetwork peering (e.g., surety of routing updates), accounting/costing for intercarrier as well as end user to carrier relationships, ensuring low latency control mechanisms, and nomadic/remote high speed access. A Public Key Infrastructure (PKI) that interacts with the industry-wide PKI is paramount to the success of integrating and deploying security in the NGI. This subtask will also develop ways for organizations or individuals to interoperate in the face of a rich and dynamic set of policies, for example, those that exist among different Federal agencies.

2.1.3.1 Cryptographic Technology and Applications

The cryptographic activity involves the development and testing of cryptographic algorithms and interfaces for use in protecting office and electronic commerce applications and data. This is one of the most important areas of information technology (IT) security, since several critical security services (authentication, data integrity, data confidentiality, and nonrepudiation) all depend on cryptographic technology.

Public Key and Key Management Infrastructure. The use of cryptographic services across the global Internet requires the use of “certificates” that bind cryptographic keys and other security information to specific users or entities in the network. Comprehensive certificate management mechanisms and underlying support infrastructure are required for all of this to work.

Internet/Internetwork Security. The security and success of the NGI will depend on the existence of new, more secure protocols. Current protocols have limited and demonstrably weak security mechanisms. New protocols being developed will include advanced methods to authenticate communications, nodes, and users, and will provide other security services such as confidentiality protection, extended audit trails, and threat monitoring.

2.1.3.2 Security Criteria, Test Methods, and Testing

Unlike other types of standards and open specifications, security standards have an implicit requirement for special testing. In addition to traditional functionality and interoperability tests, security products and services need to be tested to ensure that they cannot be subverted by intentional acts or attacks and that they do not contain functionality beyond that specified (e.g., “trap doors”). This subtask will continue efforts to develop a common set of security product/system evaluation and testing criteria to meet this need.

2.2 AGENCY SPECIFICS

2.2.1 Defense Advanced Research Projects Agency

2.2.1.1 Introduction

DARPA’s goals in network research mesh well with those of the initiative in key areas. Together with the terabit-per-second (Tbps) network speeds enabled by Goal 2.2 technologies, the network management and end-to-end QoS technologies developed under Goal 1 will be the pathway to a robust, scalable, shared infrastructure supporting lead DoD users, other government agencies, and the research community, as well as large numbers of commercial users. This will be accomplished through a partnership with industry. Specifically, the DARPA NGI program goals are as follows:

1. Develop network growth engineering and end-to-end QoS technologies and put them into the hands of the research community
2. Ensure that these technologies can be tailored for use in mission critical environments
3. Create an infrastructure partnership through which lead users (government and research) share facilities with the general public, thereby accelerating the development and penetration of novel network applications

DARPA will play a lead role in multiagency planning and in execution of NGI Goal 1 tasks with participation from DoE, NIST, NASA, and NSF. Portions of DARPA’s existing Quorum program in global distributed system technologies will form the basis for the end-to-end QoS thrust. Other agency programs will complement and leverage the Quorum program in developing advanced network services, QoS, and security technologies.

2.2.1.2 Network Growth Engineering

The goals of this task are to (1) create and deploy tools and algorithms for planning and operations that guarantee predictable end-to-end performance at scales and complexities that are 100 times those of the current Internet; (2) facilitate management of large scale internetworks operating at terabit speeds supporting a range of traffic classes on a shared infrastructure; and (3) create an infrastructure partnership through which lead users share facilities with the general public to accelerate the development and penetration of novel network applications.

This task will develop and integrate technologies for network planning and simulation; network monitoring, analysis, and control; innovative data delivery; and shared infrastructure management for lead users. The highly automated services envisioned in this task lead to the goal of building strong security mechanisms into the components.

Planning and Simulation. Planning large network interconnections is at present mainly a manual process that is not tied to any runtime tools or distributed efforts. Under this subtask, a network planning description language will be developed as a community standard. The target goals are to show that:

1. 100 small organizational networks can be planned and coordinated into an Internet
2. The plan can be validated against the requirements
3. Requirements can be maintained during the life cycle of the resulting network

Monitoring, Control, and Display. Network engineering and management requires tools for gathering and analyzing data and for issuing control commands based on the results. Current practice uses planned analysis based on protocol headers and aggregated statistics. This subtask will take on the challenge of presenting runtime analysis based on distributed communication patterns and communication speeds so high that they would overwhelm current tools. This will be accomplished by developing the following capabilities:

1. Monitoring and analytical tools and capabilities at all protocol levels for all speeds, but especially those at OC-48 and higher
2. Distributed control of all configuration parameters at distributed management stations
3. Large screen visualization of runtime data showing deviation from planning models, efficiency of distributed computing communication, and available command options for all hardware
4. Control of wavelength division multiplexing (WDM) cross domain circuit assignments for at least two optical networks

Integration. The goal of this subtask is to ensure that the network engineering tools and the high performance transmission and switching technologies work smoothly together for accomplishing end-to-end management of leading edge user requirements. The requirements will be developed with the users, and the tools will be integrated into the NGI testbed, developing 10 distributed management stations with monitoring and command interfaces to all connected equipment; half of the management stations will be in the Goal 2.1 testbed and half will be in the Goal 2.2 testbed.

The management stations will demonstrate the scalability nature of the design and its capability to adjust to varying levels of granularity of information, organizational responsibility, and active versus passive monitoring and control. Analysis and control functions will work smoothly among cooperating organizations, and the ability to implement wide area, distributed interoperation will be demonstrated as an ongoing NGI effort.

Data Delivery. The management software will work in a tightly bound interlock with new strategies for controlling data delivery in networks. Research efforts under this subtask will develop network interior nodes that combine methods previously seen as disjoint or mutually exclusive: routing and switching, best effort and priority traffic, dynamic routing and virtual circuits, greedy admission versus guaranteed delivery, and flat rate versus variable costing. In addition, schemes that use generalized or alternative addressing methods will be explored. Tools that permit network engineers to adjust the strategy trade-offs to best meet their requirements will be prototyped and tested in the high speed arena.

This subtask will develop the technology to allow lead users to share the same infrastructure as conventional users. The DoD, other government agencies, and the research community at large typically have lead user requirements for telecommunication facilities that are beyond those of the typical users by orders of magnitude. In some cases there may be sufficient aggregate capacity in the existing infrastructure to support their requirements; however, it is formatted or managed in ways that preclude coexistence among the lead and conventional users. Traditionally, wide area telecommunication service providers have addressed this problem by installing leased lines, a solution that has been extraordinarily expensive for the lead users. This subtask will investigate architectural concepts, management strategies, and operational arrangements that will facilitate the sharing of a common, wide area infrastructure.

2.2.1.3 End-to-End Quality of Service

The goals of this task are to facilitate the delivery of end-to-end ensured QoS to applications and to ensure that these technologies can be tailored for use by lead users who have demanding requirements. The strategy is to allow users to negotiate application specific trade-offs among such parameters as bandwidth, latency, precision, and reliability in order to obtain predictable performance at a known quality level. Exploiting emerging network level mechanisms is difficult, however, for they are semantically far removed from the applications they are intended to support and are accessible only through layers of software. End-to-end QoS assurance requires an approach that spans these operating system and middleware layers in order to effectively deliver network level QoS guarantees. This task will develop and demonstrate a comprehensive QoS management architecture; drill down technologies to facilitate propagation of QoS constraints across software layers; and next generation network technologies to support QoS.

Baseline QoS Architecture. The baseline QoS management architecture will provide the framework of models, languages, and protocols to permit distributed applications to specify multidimensional QoS requirements, to negotiate acceptable trade-offs and confidence levels, and to receive feedback on delivered QoS enabling adaptation. APIs supporting the propagation of QoS constraints and feedback through software layers will be developed. The baseline architecture will support a broad range of defense, government, and commercial applications by focusing on the fundamental QoS dimensions of performance, availability, precision, and soft real time. This will subsequently be extended under DARPA funding to cover mission critical properties. Specific areas to be investigated under this subtask include the following:

1. Specification and translation of application specific views of QoS into constraints on individual resources and propagation of those constraints through software layers to network services
2. Negotiation of QoS “contracts” providing applications with quantitative confidence bounds on the level of service to expect and explicit characterization of the trade offs involved
3. Monitoring and feedback technology to continually measure delivered QoS and to notify applications when QoS contracts can no longer be honored.

Drill Down Technologies. Current technologies support composition of *functionality* across system layers, but not the composition of their QoS properties. This task will develop technologies that can be used to drill down and expose interfaces to QoS and network management capabilities that are presently hidden within the individual layers. Emphasis will be on techniques appropriate for operating systems, communications libraries, and middleware services, including distributed objects. Specific topics to be addressed include the following:

1. Adaptation mechanisms, such as specialization, “on-the-fly” code generation, and dynamic module substitution
2. Efficient triggering and assurance mechanisms, including event and exception handling, guarded execution, monitoring
3. Coordinated adaptation across layers and cross layer optimization

2.2.1.4 Security

The NGI security goals are to provide the basis for implementing and enforcing appropriate security policies among organizations, users, and infrastructure components under shared control. Interoperable authentication methods are a prerequisite. NGI will meet the continuing challenge of building new services that use the network security architecture, but it will also develop ways for organizations or individuals to interoperate in the face of a rich and dynamic set of policies, for example, those that might exist among different Federal agencies. The assurance that security mechanisms are available, correct, and used will also be addressed.

DARPA's primary security activities under NGI will be in securing the network management functions, exploring the integration of security into the QoS architecture, and ensuring the secure activation of drill down mechanisms.

2.2.1.5 Milestones

<u>Network Engineering</u>

Planning and Simulation

FY2000 (4Q)	Demonstrate planning to meet requirements for 100,000 node, five protocol layer network; verification through simulation
FY2001 (2Q)	Demonstrate planning for 10 million nodes, seven protocol layers, real time simulation

Monitoring, Control, and Display

FY1999 (2Q)	Demonstrate inline monitoring of protocol headers for 2.5-gigabit (Gb) lines
FY2000 (2Q)	Implement standard API for control of network resources at all levels; compile commands based on evaluation of network conditions. Demonstrate 15-msec response capability
FY2001 (1Q)	Evaluate design for 100-gigabits-per-second (Gbps) monitoring
FY2001 (2Q)	Integrate large screen visualization with high speed network testbed; five distributed control stations

Integration

FY2000 (2Q)	Integrate five stations into OC-3 network; demonstrate 25 percent improvement in resource utilization over 3 month period
FY2001 (2Q)	Integrate five management stations into multigigabit network; demonstrate 100 percent improvement in resource utilization over 3 month period

Data Delivery

FY2000 (2Q)	Deliver router with three heterogeneous delivery strategies providing video, digital library, and shared whiteboard services; demonstrate 100 percent improvement in throughput over single strategy
FY2002 (1Q)	Demonstrate network elements with five or more strategies at > 100 percent improvement in utility of communication

Lead User Infrastructure Sharing

FY2000 (1Q)	Demonstrate lead user channel striped over 100 ordinary channels without performance loss for reliable data transmission (Transport Control Protocol—TCP).
FY2001 (2Q)	Demonstrate split second configuration among 100 lead users versus. 10,000 ordinary users for reliable data communication (TCP)

<u>Quality of Service</u>

Baseline QoS Architecture

FY1998 (2Q)	Demonstrate availability of QoS negotiation over wide area ATM networks using models allowing control of single fault-tolerance mechanism
FY1998 (4Q)	Define baseline QoS architecture

Drill Down Technologies

FY1998 (4Q)	Demonstrate operating system (OS) kernel adaptation tools and mechanisms to achieve a factor of 3 to 5 reduction in communications overhead

2.2.2 Department of Energy

2.2.2.1 Introduction

DoEís goals in network research strongly support the NGI goals. This program accentuates network research and infrastructure, not only to support higher speeds (i.e., end-to-end gigabits), but also to support more efficient and smarter use of network resources. DoE has a long and successful history of network research in high speed end system interfaces, protocols and services to support collaborative environments, and bearer service congestion flow control, and general management tools and techniques. In addition, DoE has always been an early adopter of advanced technologies and hence has the system level integration expertise required to integrate and evolve advanced technologies from other agency and DoE research efforts to enable its advanced applications.

Therefore, DoE will leverage its research expertise to support and complement Goal 1 activities pursued by DARPA and other agencies. DoEís efforts in Goal 1 will extend its work in areas relevant to the use and support of QoS capabilities at the network, operating system, and application layers, and will further its efforts in IP and ATM multicast, network management, and flow and congestion control. Goal 1 services and mechanisms will be developed in and incorporated into DoEís NGI Goal 2.1-supported 100x infrastructure as soon as possible.

Quality of Service and network management capabilities are two critical requirements for enabling DoEís advanced distributed and collaborative applications. In QoS, DoE will continue its cooperative involvement with DARPAís Quorum research activities. It will also explore both Internet Protocol based (IP) and ATM based resource and admission control, scheduling, management, prioritization, accounting (bidding, costing, etc.), authentication, analysis, monitoring, assurance, and debugging mechanisms to support both application based QoS invocation control and to support site and carrier/Internet Service Provider (ISP) administrators with their management tasks of IP, ATM, and other technology networks. DoE will also investigate how to maximize its use of these high end services and technologies through the support of network research and production traffic on as much of the same infrastructure as possible (i.e., Goal 2.1 100x networks).

DoE has identified the following as possible areas in which to support Goal 1, but will also pursue other areas that have not yet been identified, as well as to leverage DARPA and other agency programs. It is expected that initial activities will commence in 1998 and that they will evolve and be refined through cooperative efforts with other agency researchers and de facto standards such as the Internet Engineering Task Force (IETF) and ATM forum. Yearly solicitations will be made for research on these and other relevant topics, and the resulting proposals will be coordinated with other agency activities by using other agency program managers and experts as reviewers and through interagency program manager coordination.

2.2.2.2 Network Growth Engineering

Planning and Simulation. DoE will contribute its expertise on network traffic characterization, analysis, and monitoring (e.g., the Network Probe Daemon (NPD)) to DARPA efforts to simulate networks and to the NSF sponsored activities at the National Laboratory for Applied Network Research (NLANR).

Monitoring, Control, Analysis, and Display. DoE will develop and deploy 100x capable network management and monitoring tools that can be used by end users and by network administrators, and that provide assurance of QoSand support integrated and “drill down” or “cross layer” access to the network layers, debugging, analysis, and monitoring of the IP bearer service, ATM level, and other relevant technologies involved in high speed and advanced end-to-end connectivity. Advanced tools will continue to be needed to characterize, monitor, and analyze network traffic. DoE will also develop and deploy nonintrusive active and passive advanced network monitoring and service assurance agents, servers, and capabilities that can be used on NGI Goal 2.1 and Goal 2.2 networks.

Integration. DoE will lead the effort to provide effective and capability rich interagency interconnections and peering points, complete with appropriate management tools, for the Goal 2.1. network. This will include research and deployment on appropriate multilayer interconnection and peering architectures that support both production quality services to the NGI applications and network research on as much of the same wide area network (WAN) and campus infrastructure as possible.

Data Delivery. DoE will continue work on IP and ATM based congestion and flow control techniques and mechanisms. This involves continuing work on the Random Early Drop (RED) congestion technique in high speed networks and the use of caching of different types of traffic to reduce traffic congestion and bottlenecks, as well as both reliable and unreliable multicast and Real Time Protocol (RTP).

The agency will also develop and deploy tools and mechanisms to enhance our ability to interconnect both carrier and Internet Service Provider (ISP) networks so as to sustain end-to-end QoS, security, network and management, and to minimize fragmentation induced inefficiencies.

Managing Lead User Infrastructure. DoE will investigate and deploy an architecture that supports both network research and production (advanced application) traffic on as much of the same infrastructure as possible. This will enable the applications to concurrently stress the new technologies and to take advantage of the advanced preproduction capabilities without going through massive transitions.

DoE will also investigate the extension of QoS capabilities and “striped” network access into the operating system (workstation, parallel systems, storage servers), and provide for efficient application to low level service plane control and framing (e.g., direct application control of WDM frequencies or packets over the Synchronous Optical Network—SONET).

2.2.2.3 End-to-End Quality of Service

Baseline QoS Architecture. DoE will develop and deploy admission control, scheduling, management, prioritization, accounting (bidding, costing, etc.), authentication, analysis, monitoring, assurance, and debugging mechanisms to support both application based QoS invocation control and support site and carrier/ISP administrators with their management tasks of IP, ATM, and other technology networks. This will require work with IP based RSVP, as well as ATM QoS; therefore, DoE will continue its cooperative work with the DARPA Quorum program and with any others that lead to the same goal. In addition, we will explore extending QoS over last meter technologies (e.g., Gigabit Ethernet) that do not traditionally support QoS. The goal is to develop, enhance, incorporate, and integrate as many of these new technologies into the DoE NGI Goal 2.1 and 2.2 networks on an end-to-end (i.e., application to application) basis as quickly as possible.

DoE will also develop a QoS API that provides for semantic mapping of QoS from the application perspective to that provided by the underlying services and provide for cross-layer signaling and triggering of QoS mechanisms. This API will support Accelerated Strategic Computing Initiative (ASCI) and DoE 2000 applications.

Drill Down Technologies. DoE will expose network management capabilities to applications, providing direct access to network layer components and objects.

Integrated cross-layer debugging and analysis tools and techniques will also be developed. This task will complement and leverage DARPA's Quorum program.

Next Generation Network Technologies. DoE will develop an authenticated resource broker and services for class based queuing (CBQ) and link sharing and deploy CBQ into a pre-production high speed NGI network, as well as investigate its adaptation and scaling properties.

2.2.2.4 Security

Security is essential to the success of the NGI. It is necessary for supporting a secure and fair means of user access to and use of network resources (e.g., QoS), smart network management, internetwork peering (e.g., surety of routing updates), accounting/costing for intercarrier as well as end user to carrier relationships, ensuring low latency control mechanisms, and nomadic/remote high speed access. A Public Key Infrastructure (PKI) that is integrated into and interacts with the industry-wide PKI is paramount to the success of integrating and deploying security in the NGI. The following goals will be pursued:

1. Integrate authentication and security mechanisms into QoS and admission control mechanisms and into the general network NGI Goal 2.1 architecture
2. Secure the network management and monitoring functions at the campus local area network (LAN) and WAN level
3. Provide secure mechanisms to support remote online facilities and collaborative environments
4. Implement and deploy an industry interoperable PKI that can be used by DoE and its interagency collaborators and applications

2.2.2.5 Milestones

The milestones noted below may span multiple years because many of these research areas will evolve over time from base implementations to usable tools or capabilities. The beginning date identifies when the first instance of a tool or capability is expected, and the last date indicates when a fairly robust usable version is expected. Demonstration of these capabilities will be done on a yearly basis. The dates may also fluctuate as a result of the research proposals received in response to solicitations and as a result of similar work at other agencies.

<u>Network Growth Engineering</u>

Monitoring, Control, and Display

FY1998-1999	Provide ATM probe and servers for OC-3, OC-12, and up on DoE NGI Goal 2.1 sites
FY1998-1999	Integrate IP and ATM debugging, monitoring, analysis tools
FY1998-1999	Extend both scale and speed of IP-based Network Probe Daemon to support NGI measurement infrastructure
FY1998-2000	Deliver gigabit speed monitoring and analysis tool
FY1999-2000	Deliver RSVP, as well as admission control, analysis, and debugging tools
FY1998-2002	Demonstrate distributed NOC and interNOC capabilities for NGI Goal 2.1
FY1999-2002	Deliver ATM QoS/Network-to-Network Interface (NNI) analysis and debugging tools

Integration

FY1998-1999	Interconnect DoE NGI infrastructure with other agency NGI networks, for example, the Very High Speed Backbone Network Service (vBNS) and the NASA Research and Education Network (NREN)
FY1998-1999	Develop cross-carrier interconnect and multiinstitution peering management and analysis tools

Data Delivery

FY1998-1999	Extend RED to high speed routers, switches, and networks.
FY1998-2000	Deliver striping of very high speed data and control channels into end systems
FY1998-2000	Continue development of RTP and multicast technologies and analysis tools

Lead User Infrastructure Sharing

FY1998	Accomplish interagency/intercarrier interconnection by means of agency equipment (IPv4)
FY1999-2000	Provide interagency/intercarrier QoS support for ATM (NNI)
FY1998-2001	Direct intercarrier/interagency NGI ATM and IPv6 interconnection
FY1999-2001	Deliver dual-mode (production and R&D) end node
FY1998-2002	Demonstrate 6bone (IPv6 backbone) participation via the Energy Science Network (ESnet)
FY1999-2002	Investigate (de)aggregating and striping of campus and WAN infrastructure, (SONET, WDM, for example, High Performance Parallel Interface 64 (HIPPI64), and ATM)
FY1998-2002	Provide advanced multimode infrastructure (i.e., concurrently support production and network research) on campus and WAN
FY1999-2002	Provide interagency/intercarrier QoS support for RSVP

<u>End-to-end QoS</u>

Baseline QoS Architecture

FY1998-2000	Deliver QoS API for DoE 2000, ASCI Grand Challenge applications
FY1998-2002	Provide authentication (including PKI) for QoS
FY1998-2002	Deliver network viewing and control capabilities for applications and libraries

Drill Down Technologies

FY1998-1999	Demonstrate integrated cross-layer analysis and debugging tool
FY1998-2000	Provide application drill down access to various network layer planes

Next Generation Network Technology

FY1998-2000	Demonstrate cross-layer QoS signaling
FY1999-2000	Demonstrate ClassBased Queuing (CBQ) with relevant resource broker, admission control, and analysis and debugging tools
FY1999-2000	Investigate QoS support over last meter non-QoS media
FY1998-2001	Provide admission control/cost accounting, etc., for RSVP
FY1998-2002	Provide I-NNI and P-NNI QoS support for ATM
FY1998-2002	Provide Admission control/cost accounting for ATM
FY1998-2002	Enable 6bone participation to develop and analyze IPv6 capabilities
FY2000-2002	Provide native deployment of IPv6 on NGI

Security

FY1998-2000	Develop DoE-wide certificate hierarchy and PKI that is interoperable with industry
FY1998-2000	Secure high speed and latency bounded access to DoE online facilities
FY1998-2002	Provide authentication (including PKI) for QoS, admission control, accounting/costing, etc
FY2000-2002	Secure network routing updates on NGI Goal 2.1 network

2.2.3 National Institute of Standards and Technology

2.2.3.1 Introduction

NIST is refocusing part of its on-going research program in advanced networking technologies, computer security, and conformance testing to better support NGI goals. The NGI initiative provides an exceptional opportunity for NIST to coordinate its on-going research with other Federal agencies, to build on the investments of those agencies, and to refine its research focus based on the results of their programs of work. NIST's programs focus on measurement, standards, and test methods to expedite the development of, improve the quality of, and enhance the interoperability of next generation networking technologies and supporting systems. NIST's unique mission in these activities is to support the U.S. information technology industry by fostering the rapid commercialization and deployment of enabling and infrastructural networking technologies developed as part of the NGI effort.

NIST's traditional focus on measurement, standards, and test methods in support of the IT industry will take on new importance as the research efforts of Goal 1 evolve toward standardized/commercial technologies. The complexities of, and interdependencies among, future network control systems and services (e.g., multilayer QoS signaling, routing, flow control, security) will defy analysis by simple means. Likewise, the cost of building testbeds of the extent necessary to test many critical design and deployment issues (e.g., global scaling, highly layered control systems) will be prohibitive in the early stages of technology R&D. The old credo of "rough consensus and running code" being the only quality assurances needed to guide the development of new Internet technology will not be sufficient in the future.

NIST's goals are to research and develop new techniques and tools to test and evaluate new networking technology at all stages of its development and deployment. Test and instrumentation technology should become a common part of the protocol design and specification process, and should be integrated into, and evolve with, the implementation and deployment of the network itself.

Advances in measurement and testing technology will enable the rapid evaluation of research designs and prototypes and will facilitate the transfer of new technology to the communications industry and the NGI Goal 2 network infrastructures.

NIST has initially identified and refocused as necessary the following on-going research activities that support Goal 1. It is expected that future NGI activities will evolve from these efforts and general goals. Participating in the NGI initiative will give NIST a broader basis for evaluating the effect of its research and the direction that future research should take. Future research objectives will be refined through collaboration and cooperative efforts with the NGI programs of other agencies and with the needs of the IT industry in adopting and commercializing NGI technology.

Next Generation Internetwork Technology

NIST activities in support of fostering commercialization of next generation networking technologies include developing a reference implementation and a remotely accessible interoperability testbed for IPv6; developing an Integrated Services Packet Switched network testbed and instrumentation tools for QoS in IP-based networks; simulating and analyzing protocols for advanced ATM networks; doing research in high speed residential access; and testing and measurement methods for evaluating next generation, intelligent collaborative tools.

IPv6 Technology Development. NIST is actively participating in the design, development, and testing of the NGI protocol, IP version 6 (IPv6). NIST's activities include (1) pilot deployment of a multivendor IPv6 testbed focusing on evaluating IPv6 security mechanisms and technical migration strategies for existing IP networking infrastructures; and (2) the development of publicly available reference implementations of IPv6 focusing on security features and ATM integration issues.

Integrated Services for NGI. NIST is participating in the Internet Engineering Task Force (IETF) community effort to develop technology to support real time network services in the Internet Protocol Suite. NIST's activities include (1) pilot deployment of a multivendor testbed focusing on evaluating the interoperability of early prototype implementations of resource reservation protocols and real time transport protocols, and (2) the development of instrumentation and emulation tools that enable application experimentation with emerging QoS controlled network services.

ATM Protocol Simulation and Analysis. NIST has built the ATM protocol simulator that is widely used within the ATM Forum to analyze proposed ATM protocols. NIST uses the simulator to analyze a wide range of ATM protocols. Current examples of this work are the proposed traffic management and ATM network routing protocols. This simulation and analysis capability can be used for modeling and evaluating technical proposals for mapping IP reservations to ATM QoS or for analyzing various proposals for IP switching on an ATM fabric.

Research in High speed Residential Access. NIST is performing the media access layer analysis for the proposed Hybrid Fiber Coax (HFC) standard for the IEEE 802.14 committee. HFC allows high speed two-way communication through a home's cable TV connection. Handling the aggregate demands of large numbers of high speed residential customers may be a bigger challenge for the NGI infrastructures of Goal 2 than handling individual very high speed applications. Traffic characterization models of residential customers and others, and models of the effect of aggregate demand are important test and measurement tools for network planning.

Testing Technology for Collaborative Systems. NIST is developing method, metrics, and testing tools to evaluate generation-after-next collaboration systems and supporting infrastructure technologies. NIST's efforts focus on developing testing and instrumentation technology that will enable collaboration systems developed through the DARPA Intelligent Collaboration and Visualization (IC&V) program to be evaluated in terms of objectives for task performance, scalability, heterogeneity, and interoperability.

2.2.3.3 Security

NIST activities in support of secure systems and networks include the development of criteria, tests, and test methods for Internet/internetwork security, cryptographic technology, advanced authentication technology, and public key and key management infrastructure.

Cryptographic Technology and Applications. This activity involves the developing and testing cryptographic algorithms and interfaces for use in protecting office and electronic commerce applications and data. This is one of the most important areas of IT security, since several critical security services (authentication, data integrity, data confidentiality, and non-repudiation) depend on cryptographic technology.

Advanced Authentication Technology. Reusable passwords remain the primary means of user authentication in the Internet, despite the fact that it has been demonstrated that their use is completely unsuitable for the global, open Internet where such passwords are passed in the clear and can be picked up and reused almost at will. NIST has led in the development of technology and standards to provide effective authentication alternatives to passwords, and is working to promote the use of such methods throughout the Internet and other network environments.

Public Key and Key Management Infrastructure. The use of cryptographic services across the global Internet requires the use of “certificates” that bind cryptographic keys and other security information to specific users or entities in the network. Comprehensive certificate management mechanisms and underlying support infrastructure are required for all of this to work. NIST is actively involved in developing critical components of that public key infrastructure (PKI).

Internet/Internetwork Security. The security and success of the NGI will depend on the existence of new, more secure protocols. Current protocols have limited and demonstrably weak security mechanisms. New protocols being developed (by the Internet community in general, with the active involvement of NIST) will include advanced methods to authenticate communications, nodes, and users, as well as provide other security services such as confidentiality protection, extended audit trails, and threat monitoring. Remote access to the NIST/National Security Agency (NSA) IP security reference implementation will be developed for remote interoperability testing. Web-based interoperability testing will also be developed.

Security Criteria, Test Methods, and Testing. Unlike other types of standards and open specifications, security standards have an implicit requirement for special testing. In addition to traditional functionality and interoperability tests, security products and services need to be tested to ensure that they cannot be subverted by intentional acts or attacks and that they do not contain functionality beyond that specified (e.g., “trap doors”). NIST has helped lead an international effort to develop a common set of security product/system evaluation and testing criteria to meet this need. This is central to NIST’s strategy of developing (1) a comprehensive testing competency, (2) research, and (3) accreditation capability for use by product developers and users in both government and industry.

2.2.3.4 Milestones

FY1998-2001	Demonstrate IP integrated services interoperability testbed
FY1998-2001	Demonstrate IPv6/IPSec interoperability testbed
FY1998-2001	Provide prototype network emulation tool for QoS-sensitive applications
FY1998-2002	Provide prototype IP-integrated services protocol instrumentation tool
FY1998-2002	Demonstrate IPv6/IPSec prototype
FY1998-2002	Develop test and evaluation methods for collaborative applications
FY1999-2002	Demonstrate full PKI
FY1999-2002	Validate tools and techniques for testing collaborative applications
FY1998-2003	Demonstrate IPv6/ATM prototype
FY1998-2003	Demonstrate Web-based Internet Protocol (Secure) (IPSec) interoperability testing tool

FY1998-2003	Demonstrate IPv6/IPSec/ISAKMP prototype
FY1998-2004	Demonstrate IPv6 PKI
FY1998-2004	Provide security criteria and tests

2.2.4 National Aeronautics and Space Administration

2.2.4.1 Introduction

NASA's goals in network research support NGI goals in key areas by balancing networking research and applications networking through the increased functionality of Goal 1. NASA will continue to be an early adopter of emerging networking technologies that chart a course for a robust, scalable, shared infrastructure supporting lead users from NASA, other government agencies, and the research community, as well as large numbers of commercial users.

NASA's program goal relevant to NGI Goal 1 is to sponsor research and development (R&D) in new networking technologies and services in support of the high performance applications requirements. By partnering with industry and academia on R&D in internetworking technologies to achieve an interoperable high performance network testbed, NASA will deliver advanced networking technologies to the aerospace community and ultimately to the public.

Specifically, the NASA Research and Education Network (NREN) project and its existing network will provide a basis for implementing the NASA NGI plan. The NASA NGI program goals are as follows:

1. Introduce next generation internetworking technologies into NASA mission applications
2. Create an infrastructure partnership through which lead users (government and research) share facilities with the general public, thereby accelerating the development and penetration of novel network applications
3. Ensure that technologies are transferable, and that they integrate and scale properly to production networks

NASA will collaborate with DoE, NIST, DARPA, and NSF in planning and executing Goal 1. NASA will deploy an appropriate suite of advanced networking services to enable high performance applications. NASA sponsored research will focus on issues such as network performance measurement, network interoperability scaling, management, QoS, and network security. NASA will fund and manage research in advanced network technologies that are richer in features, higher in performance, and deliverable at a reasonable cost. For example, they will enable real time networking, group collaborations, and a seamless interface for space-to-ground communications.

2.2.4.2 Network Growth Engineering

The goals of this task are to (1) create and deploy tools and algorithms for planning and operations that guarantee predictable end-to-end performance at scales and complexities 100 times those of the current Internet; (2) facilitate management of large scale internetworks operating at terabit speeds supporting a range of traffic classes on a shared infrastructure; and (3) create an infrastructure partnership through which lead users (government and research) share facilities with the general public, thereby accelerating the development and penetration of novel network applications. This task will develop and integrate technologies for network planning and simulation; for network monitoring, analysis, and control; for innovative data delivery; and for shared infrastructure management for lead users. The highly automated services envisioned in this task lead to the goal of building strong security mechanisms into the components.

Planning and Simulation. Planning large network interconnections is now primarily a manual process that is not tied to any runtime tools or distributed efforts. NASA will utilize network modeling and analysis tools to simulate proposed networks and to develop models that can be used as baselines for early designs. The target goals are to show the following:

1. That 100 small organizational networks can be planned and coordinated into an Internet
2. That the plan can be validated against the requirements
3. That requirements can be maintained during the life cycle of the resulting network.

Monitoring, Control, and Display. NASA will collaborate with DoE and DARPA on the deployment of 100x-capable network management and monitoring tools, which can be used by end users and by network administrators, that provide assurance of QoS and that support integrated and “drill down” or “cross layer” access to the network layers, debugging, analysis, and monitoring of the IP bearer service, ATM-level, and other relevant technologies involved in high speed and advanced end-to-end connectivity. Advanced tools will continue to be needed to characterize, monitor, and analyze network traffic.

NASA will also deploy nonintrusive active and passive advanced network monitoring and service assurance agents, servers, and capabilities that can be used on NGI Goals 2.1 and 2.2 networks.

Integration. NASA, DoE and NSF will lead the effort to provide effective and capability-rich interagency interconnections and peering points, complete with appropriate management tools, for the Goal 2.1 network. This will include research on and deployment of appropriate multilayer interconnection and peering architectures that support both production quality services to the NGI applications and network research on as much of the same WAN and campus infrastructure as possible.

Data Delivery. NASA will support experiments in new strategies for controlling data delivery in networks. NASA will support DARPA's research efforts under this subtask to develop network interior nodes that combine methods previously seen as disjoint or mutually exclusive: routing and switching, best effort and priority traffic, dynamic routing and virtual circuits, greedy admission versus guaranteed delivery, and flat-rate versus variable costing. In addition, schemes that use generalized or alternative addressing methods will be explored. Tools that permit network engineers to adjust the strategy trade-offs to best meet their requirements will be prototyped and tested in the high speed arena.

Managing Lead User Infrastructure. This subtask will develop the technology to allow lead users to share the same infrastructure as conventional users. NASA, other government agencies, and the research community at large typically have lead user requirements for telecommunication facilities that are orders of magnitude beyond those of the typical users. In some cases there may be sufficient aggregate capacity in the existing infrastructure to support their requirements, but it is formatted or managed in ways that preclude coexistence among the lead and conventional users. Traditionally, wide area telecommunication service providers have addressed this problem by installing leased lines, a solution that has been extraordinarily expensive for the lead users. This subtask will investigate architectural concepts, management strategies, and operational arrangements that will facilitate the sharing of a common, wide area infrastructure.

2.2.4.3 End-to-End Quality of Service

The goals of this task are to facilitate the delivery of end-to-end ensured QoS to applications, and to ensure that these technologies can be tailored for use by lead users who have demanding requirements. The strategy is to allow users to negotiate application specific trade-offs among such parameters as bandwidth, latency, precision, and reliability in order to obtain predictable performance at a known quality level. Exploiting emerging network level mechanisms is difficult, however, for they are semantically far removed from the applications they are intended to support and are accessible only through layers of software. End-to-end QoS assurance requires an approach that spans these operating system and middleware layers in order to effectively deliver network level QoS guarantees. This task will develop and demonstrate a comprehensive QoS management architecture, drill down technologies to facilitate propagation of QoS constraints across software layers, and next generation network technologies to support QoS.

Baseline Quality of Service Architecture. The baseline QoS management architecture will provide the framework of models, languages, and protocols to permit distributed applications to specify multidimensional QoS requirements, negotiate acceptable trade-offs and confidence levels, and receive feedback on delivered QoS enabling adaptation.

This effort will deploy admission control, scheduling, management, prioritization, accounting (bidding, costing, etc.), authentication, analysis, monitoring, assurance, and debugging mechanisms to support both application based QoS invocation control and support site and carrier/ISP administrators with their management tasks of IP, ATM, and other technology networks. This will require work with IP-based RSVP and ATM QoS; therefore, NASA will collaborate with the DARPA Quorum program and with any others that lead to the same goal. The goal is to develop, enhance, incorporate, and integrate as many of these new technologies into the NASA portion of the NGI Goal 2.1 and 2.2 networks on an end-to-end (i.e., application-to-application) basis as quickly as possible. NASA will leverage and complement its efforts in advocacy of advanced protocol development and enhancements related to space-terrestrial communications (e.g., IPv4/IPv6 and ATM APIs) with respect to satellite links, and will continue its involvement and leadership in standards development through the various standards bodies and organizations (IEEE, ATM Forum, etc.).

NASA will also develop a QoS API that provides for semantic mapping of QoS from the application perspective to that provided by the underlying services and provide for cross-layer signaling and triggering of QoS mechanisms. This API will support High Performance Computing and Communications (HPCC) and Mission to Planet Earth (MTPE) applications.

Drill down Technologies. Current technologies support composition of *functionality* across system layers, but not the composition of their QoS properties. NASA will focus on exposing network management capabilities to applications, and on providing direct access to network layer components and objects. Integrated cross-layer debugging and analysis tools and techniques will also be deployed in partnership with DoE. This task will complement and leverage DARPA's Quorum program.

2.2.4.4 Security

The NGI security goals are to provide the basis for implementing and enforcing appropriate security policies among organizations, users, and infrastructure components under shared control. Interoperable authentication methods are a prerequisite. NGI will meet the continuing challenge of building new services that use the network security architecture, but it will also develop ways for organizations or individuals to interoperate in the face of a rich and dynamic set of policies, for example, those that might exist among different Federal agencies. The assurance that security mechanisms are available, are correct, and are used will also be addressed.

NASA's primary security activities under NGI will be in securing the network management functions, exploring the integration of security into the QoS architecture, and ensuring the secure activation of drill down mechanisms.

2.2.4.5 Milestones

<u>Network Growth Engineering</u>

Planning and Simulation

FY1998 (3Q) Develop standard simulation models to “grow” internetwork/intranetwork and develop baseline simulation statistics

FY1999 (1Q) Utilize requirement analysis and configuration management procedures to design and manage the NASA virtual testbed

Monitoring, Control, and Display

FY1999 (2Q) Develop distributed NOC and inter-NOC capabilities for NGI Goal 2.1

FY1999 (2Q) Provide integrated IP and ATM debugging, monitoring, and analysis tools

FY1999 (3Q) Deploy ATM probe and servers for OC-3, OC-12, and up on DoE NGI Goal 2.1 sites

FY2000 (2Q) Provide RSVP, admission control, analysis and debugging tools

FY2000 (3Q) Deliver gigabit-speed monitoring and analysis tools

FY2000 4Q) Develop solutions for network monitoring and management tools for maintaining and measuring performance on NASA testbeds

FY2000 (4Q) Test network and transport protocols, encryption, and network management tools for high performance network

FY2001 (3Q) Demonstrate remote network configuration and control at five sites across NASA testbeds

FY2001 Demonstrate network viewing and control capabilities for applications

FY2002 (3Q) Provide ATM QoS/NNI analysis and debugging tools

FY2002 (3Q) Demonstrate distributed network management and monitor tools across five NASA testbeds

Integration

FY1998 (2Q) Demonstrate interconnection of NASA NGI infrastructure with other agency NGI networks (e.g., NSF vBNS, DoE NGI)

FY1998 (3Q) Develop plan with satellite community to collaborate on network service enhancements for achieving end-to-end seamless interoperability across high speed terrestrial/satellite network links

FY1999 (4Q) Develop cross-carrier interconnect and multi-institution peering management and analysis tools

FY2002 (4Q) Demonstrate distributed management tools that cross multiple organization and vendor networks to meet agreed service levels and to ensure interoperability

Data Delivery

FY1998 (2Q) Implement native multicast protocols on three NASA testbeds

FY1999 (2Q) Demonstrate IP support for multimedia and real time audio and video across NASA testbeds

FY2000 (3Q)	Demonstrate IP support for multimedia and real time audio and video across NGI
FY2000 (4Q)	Develop multicast as a reliable service with acknowledged delivery and authentication. Demonstrate reliable multicast on five NASA testbeds

Lead User Infrastructure Sharing

FY1998 (3Q)	Demonstrate interagency/intercarrier interconnection by means of agency equipment (IPv4)
FY1999 (2Q)	Implement dual mode (network research and application network) across network
FY2000 (1Q)	Demonstrate interagency/intercarrier QoS support for ATM (NNI)
FY2000 (2Q)	Determine (de)aggregation schemes across campus and WAN infrastructure (e.g., SONET, WDM, and ATM)
FY2001 (2Q)	Direct intercarrier/interagency NGI ATM and IPv6 interconnection
FY2002 (2Q)	Demonstrate interagency/intercarrier QoS support for RSVP

<u>Quality of Service</u>

Baseline QoS Architecture

FY1998 (3Q)	Characterize QoS and analyze requirement of multimedia protocols (e.g., MPEG-2 over ATM)
FY1998 (4Q)	Implement QoS parameters on five NASA testbed sites
FY1999 (4Q)	Implement resource reservation and real-time protocols on five NASA testbed sites
FY2001 (3Q)	Demonstrate guaranteed bandwidth and network availability on five sites across NASA testbeds
FY2001 (4Q)	Demonstrate network reliability, QoS, scalability, bandwidth-sharing and integrated network services across NASA testbeds

Drill Down Technologies

FY1999 (4Q)	Demonstrate integrated cross-layer analysis and debugging tool
FY2000 (4Q)	Demonstrate application drill down access to various network layer planes

Next Generation Network Technology

FY2000 (4Q)	Demonstrate native deployment of IPv6 on NGI
FY2001 (2Q)	Demonstrate admission control/cost accounting, etc., for IP QoS
FY2002 (3Q)	Demonstrate admission control/cost accounting for ATM

<u>Security</u>

FY1999 (1Q)	Demonstrate authentication (including PKI) for QoS, admission control, accounting/costing, etc.

FY1999 (4Q)	Develop solutions for network security using encryption across NASA testbeds
FY2000 (3Q)	Demonstrate secure high speed and latency-bounded access to NASA online facilities
FY2000 (4Q)	Demonstrate NASA-wide certificate hierarchy and PKI that is interoperable with industry

2.2.5 National Science Foundation

2.2.5.1 Introduction

NSF projects make a central contribution to NGI goals by leveraging extensive campus and industry partnerships to connect about 100 leading research universities with a high performance network fabric, interconnecting this fabric with that of other Federal and foreign research networks, implementing and testing advanced technologies, and supporting hundreds of advanced applications. The present goals of the NSF programs for high performance connections reflect all aspects of the NGI: high performance connections to about 100 universities and their research partners with QoS technology supporting hundreds of meritorious applications.

NSF has an experienced and capable research base on which to build its existing high performance connections to the NSF supercomputer centers and their partners (Partnerships for Advanced Computational Infrastructure, PACI), the NSF-sponsored National Laboratory for Applied Networking Research (NLANR), and dozens of funded individual investigators in university and industry laboratories, as well as ongoing funded research with investigators in DARPA, DoE, NASA, and other agencies. Moreover, NSF's partnership with Internet 2 will focus the collective efforts of over 100 leading universities on next generation network technologies and related issues of deployment, management, and testing for NGI Goal 1.

2.2.5.2 Network Growth Engineering

The goals of this task are to (1) create and deploy tools and algorithms for planning and operations that guarantee predictable end-to-end performance at scales and complexities of 100 times those of the current Internet; (2) facilitate management of large scale internetworks operating at terabit speeds supporting a range of traffic classes on a shared infrastructure; and (3) create an infrastructure partnership in which lead users (government and research) share facilities with the general public, thereby accelerating the development and penetration of novel network applications.

Planning and Simulation. NSF will work with approximately 100 leading universities and the Internet 2 organization on policies and technologies for planning large scale, high performance networks. Internet 2 and its members will be coordinating 15 to 20 gigaPOPS linking institutional, state, and regional high performance networks into a national system that must support a wide variety of new technologies and applications.

Monitoring, Control, and Display. NSF and its National Laboratory for Applied Networking Research are leaders in the deployment of network measuring and monitoring devices and in the theoretical analysis of network behaviors. NSF is also funding related measurement efforts by Merit, the Common Solutions Group, and others. All will be coordinated to focus on the study of high performance networks.

Integration. NSF will work with DARPA, NASA, and DoE to support the integration of agency research networks and the sharing, where possible, of leading edge research and robust production infrastructure. Integration will be a major activity and interest at the 15 to 20 Internet 2 gigaPOPs.

Data Delivery. NSF's NLANR will extend its online, real time tool set for the visualization of network configuration and performance, focusing on the complexities of integrated networks that are of mixed technologies.

Managing Lead user Infrastructure. Lead user management strategies and technologies will be a central focus of the NSF Partnerships for Advanced Computing Infrastructure (PACI). Both the San Diego Supercomputing Center (SDSC) and the National Center for Supercomputing Applications (NCSA) science will depend heavily on successful approaches to these problems, and both PACI and the Internet 2 universities will devote considerable expertise and effort (over 200 lead network managers and technologists) toward practical, common solutions.

2.2.5.3 End-to-End Quality of Service

The goals of this task are to facilitate the delivery of end-to-end assured QoS to applications and to ensure that these technologies can be tailored for use by lead users who have demanding requirements.

Baseline Quality of Service Architecture. The baseline QoS management architecture will provide the framework of models, languages, and protocols to permit distributed applications to specify multidimensional QoS requirements, to negotiate acceptable trade-offs and confidence levels, and to receive feedback on delivered QoS enabling adaptation. NSF and its Internet 2 partners will select, test, and deploy technologies for QoS on the vBNS and connected campus networks at an accelerated schedule. QoS related activities are perhaps the major technology goal of Internet 2, which will be exploring multiple solutions and their interoperation at 15 to 20 gigaPOPs (Points of Presence) and at about 100 campuses in addition to the national networks. NSF will coordinate these activities closely with DARPA, DoE, and NASA to help achieve an early, interoperable implementation.

Drill down Technologies. NSF's vBNS presently supports simple drill down technologies such as cut through routing. This early effort will be dramatically expanded to encompass multiple approaches at the Internet 2 gigaPOPs and regional networks, and will be broadened to include the development of corresponding measurement and management tools.

Next Generation Network Technologies. NSF's Internet 2 and vBNS partners are committed to early adoption of next generation technologies such as IPv6, RSVP. These technologies will be installed and operated directly on the vBNS and other Internet 2 networks as soon as testing and management issues permit. NSF will extend the existing NLANR to include the direct participation of about 100 Internet 2 members, and will work cooperatively with them as well as with its supercomputer center sites on national scale implementation and testing of new technologies and services such as the following:

1. QoS trials in IP and ATM
2. Performance measurement and statistics
3. Caching, multicast, virtual testbeds
4. End-to-end performance tuning for systems and applications
5. Support role in universities for applications development

In addition, NSF will work with other agencies through joint solicitations to support individual R&D in network technologies including: QoS, Multicast, Ipv6, scalability, security, and others.

2.2.5.4 Security

The NGI security goals are to provide the basis for implementing and enforcing appropriate security policies among organizations, users, and infrastructure components under shared control.

NSF's primary activities related to security will consist of early implementation and testing of security technologies and policies selected in cooperation with its Internet 2 and vBNS partners. NSF will coordinate these activities for interoperability with NSA and the other Federal agencies.

2.2.5.5 Milestones

FY1998 (1Q)	Extend the NLANR to include technical staff of all connected Internet 2 sites
FY1998 (2Q)	Make joint announcements with other agencies for research awards in network technologies
FY1998 (3Q)	Provide integrated IP and ATM monitoring and analysis tools at all sites
FY1998 (3Q)	Play a formal NLANR support role in universities for applications development
FY1998 (4Q)	Implement caching and native multicast at all sites
FY1998 (4Q)	Make awards to campus PIs for new research projects
FY1999 (1Q)	Implement RSVP for weighted early drop or other QoS, etc.

FY1999 (3Q)	Implement security at all sites
FY1999 (4Q)	Implement more complex RSVP and additional services

3. GOAL 2: NEXT GENERATION NETWORK FABRIC

3.1 GOAL 2.1: HIGH PERFORMANCE CONNECTIVITY

3.1.1 Introduction

NGI Goal 2.1 will perform R&D on advanced interconnection technology to develop and deploy a wide area demonstration network fabric that will support network research and advanced network capable applications as defined in NGI Goals 1 and 3. This fabric will function as a distributed laboratory delivering a minimum of 100 times the current Internet performance on an end-to-end basis (typically greater than 100+ Mbps end-to-end) to at least 100 interconnected NGI universities, national laboratories, and Federal research sites that demonstrate research and applications that require such an infrastructure. This demonstration network fabric will be large enough to provide a full system, proof-of-concept testbed for hardware, software, protocols, and network management that is required in the commercial NGI.

It should be noted that the emphasis of Goal 2 is on terrestrial networking technologies and connectivity. Wireless technology is becoming an important element of tomorrow's Internet architecture. However, because of the intense competition in the wireless industry and insufficient government resources, NGI will highly leverage industry developments and any ongoing Federal research to provide a hybrid networking demonstration platform.

3.1.2 Strategy and Subgoals

To implement this goal, strategy and metrics will be defined for each of the following subgoals:

Infrastructure
Common bearer services
Applications feedback
Interconnections
Site selection
Network management
Information distribution and training

3.1.2.1 Infrastructure Subgoal

The large scale development of advanced applications calls for a network that is relatively stable in order to provide a consistent environment to which developers can build, and one that is also leading edge in order to incorporate new technologies as soon as they are available. Development and testing of the new network technologies requires a more flexible network, one that can be modified and tested frequently.

These objectives will be met in NGI Goal 2.1 by expanding and interconnecting existing Federal research networks into a large “leading edge but stable” network of networks for about 100 sites, augmented by a logically separate but connected network that can be configured at will to form “virtual networks” to test specific technologies and projects. This is envisioned as an advanced network infrastructure that supports the research identified in Goal 1, but that will also strive to provide a stable network to the applications. However, there will be occasions when the network may suffer partial degradation of services to support these sometimes divergent goals.

A balance must be achieved between a stable “production” environment and one that allows for research and new technology insertion so as to enable both. For the most part this can be accomplished with existing technology. It is envisioned that: (1) a relatively stable but leading edge infrastructure will be developed that will have some downtime as new technologies are developed and introduced; (2) sites will have backup commodity service; (3) a service definition will exist between the end user and providers— this will include such items as latency, demand, and scheduled service.

The strategy to form the leading edge but stable network is to connect the university community (including, for example, leading research universities and Internet 2 campuses) with the vBNS national backbone, connect the Federal laboratories and other research sites with agency research networks, and interconnect them at high performance exchanges. Selected laboratories and universities will also be connected directly to agency networks and to the new virtual network, resulting in a new fabric that can cost effectively support all of the research objectives of the NGI.

These virtual backbone networks will also seek to provide a stable infrastructure, but may potentially be less stable as they aggressively implement the research, development, and integration of new technologies and applications being developed in Goals 1 and 3. This strategy allows for sharing the costly underlying carrier infrastructure while providing a network in which the new 100x technologies can be deployed in a limited but real world environment for testing before being introduced into the full 100x testbed. These virtual backbone networks will be richly interconnected to the agency networks, as well as the 100x vBNS, gigaPOPs, and other advanced peering points as the requirements dictate.

The performance metrics that will be used to determine if the infrastructure subgoal has been achieved are the following:

1. Does the network infrastructure perform to expectations, that is, end-to-end 100x the present Internet speed?

2. Can users get the required bandwidth and capabilities that they cannot get anywhere else?
3. Does the network infrastructure accommodate Goal 1 research needs?
4. Can applications do what they need to do?
5. Has the connectivity to sites met user requirements?
6. Do the sites have access to a 100x infrastructure that supports Goals 1 and 3?

3.1.2.2 Common Bearer Services Subgoal

In order to create a relatively stable infrastructure that can be quickly deployed, IPv4 with “best effort” services will be the initial and minimum common bearer service for all agency network infrastructures. IPv6 and other advanced technologies for QoS, etc., will be deployed on the networks as soon as they demonstrate a “leading edge, but stable” performance on the virtual networks.

Direct connections in ATM, IP over SONET, and other services will be introduced on the basis of feasibility, support for Goal 1 research activities, and application demand. The use of virtual backbones over the existing agency networks to perform the required R&D will assist in developing these technologies and services, provide a relatively leading edge but stable infrastructure for applications development, and assist in solving problems of interagency and intercarrier connectivity at the ATM level.

The metrics to determine success of this subgoal will be to gather statistics on a number of connected sites, and to keep track of trends of the bearer services in use and the number of sites that transition to advanced services, that is, from IPv4 to IPv6 and ATM.

3.1.2.3 Application Feedback Subgoal

One of the basic failures of previous high end testbeds was that although the infrastructure may have been a success, the initiative did not attain its potential because there were virtually no applications that could effectively use the infrastructure. This subgoal will help to ensure that this does not occur in the NGI initiative. Not only will NGI site selection be based on the availability of NGI class applications, but it will also include application developers working with network researchers to ensure that the resulting infrastructure satisfies the application requirements and enables those applications not yet conceived. Applications developers will be intimately involved in the R&D and deployment of the NGI.

The strategy that will be used to facilitate applications development and migration onto the 100x testbed is to provide feedback from application developers to network researchers, operators, and implementors by means of regularly scheduled coordination meetings and conferences among these groups. It is envisioned that coordination meetings will occur at least four times per year over the life of the NGI initiative. The coordination conference responsibility will rotate among the agencies involved in the Goal 2.1 testbeds. All coordination conferences will be broadcast on the multicast backbone (MBone) to obtain maximum information dissemination. In addition, common links to up-to-date Web pages will be maintained by the networks and agencies involved. At a minimum, these pages shall contain information on the various applications under development on the 100x testbed, current status of the applications, problems encountered and solutions, and action items from the various meetings and conferences.

In addition, conferences that address specific issues or topics will be held on an as needed basis. The responsibility of holding the conferences will fall to the agency or group concerned with the particular issue.

In addition, to continue successful transition from the Federal testbed sector to the private sector and to ensure that applications developers have continued infrastructure support, a transition coordination group will exist for some time after the conclusion of the NGI.

The metrics that will be used to determine if this subgoal has been met are:

1. satisfaction of advanced applications requirements; and
2. the number of applications using advanced services.

3.1.2.4 Interconnection Subgoal

In order for the NGI initiative to lay the foundation for the future Internet, it is imperative that the Federal networks provide end-to-end QoS, interconnecting with advanced services at the IP and lower layers as appropriate. It is also recognized that different levels of QoS will exist as dictated by the needs of the applications, for example, email does not require the same QoS as does a real time collaborative environment used to control an experiment at a remote site.

Implementation of NGI Goal 2.1 will create a nationwide seamless NGI fabric that provides appropriate levels of QoS, as well as necessary internetwork management security. This seamless infrastructure will consist of interconnections to the Federal agency networks (i.e., the four major networks, vBNS, ESnet, NREN, Defense Research and Engineering Network–DREN) as well as the virtual NGI network architecture (described above) that support multiple network research and advanced applications views. Each agency will leverage its existing virtual domains (i.e., administrative domains with different centers of control). The agencies have already involved the major carriers. One goal of the NGI is for these carriers to cooperate seamlessly on interconnection of the Federal agency networks. To facilitate this, the agencies will form an open interagency carrier and switch vendor working group to facilitate interagency and intercarrier cooperation at the implementor level. This group will also provide input to the agencies to help define the phases for seamless interconnection across domains (i.e., IP, Permanent Virtual Circuit–PVCs, etc.) as well as transition from carrier-based to interconnection-based services.

The metrics that will be used to determine success are as follows:

1. Ability to change networks and carriers while remaining transparent to applications

2. Number of applications that successfully operate among nodes supported by at least two different carriers
3. Number of applications which successfully operate among nodes
4. A network is supported by at least two different carriers
5. Statistics and trends gathered on bandwidth and interconnects that show the migration
6. Number of users/customers who indicate they became aware of and engaged with NGI as a result of NGI information distribution or training activities

3.1.2.5 Site Selection Subgoal

Since NGI research crosses sites and agencies, a coordinated mechanism is needed to ensure that the sites have valid applications and the required infrastructure to promote the NGI Goal 2.1 success. To achieve this, each agency will identify its NGI sites and will work in a coordinated manner with the other agencies and the application group in the site selection process. The technical engineering and design issues will be addressed by the existing Federal Networking Council (FNC) Engineering and Operations Working Group (FNC-EOWG) and a Joint Federal Engineering Team (JET). The sites with valid requirements, that is, Goal 1 and NGI applications research, will present their proposals to the respective agencies that will determine the merits, based on reviews, and award the required funding or provide the required services or access.

The agencies will be responsible for ensuring that a phasing plan exists in the proposal, where appropriate, for the site to migrate to production services with non-NGI funding in the near future. They must also ensure that connection to gigaPOPs is present where required; that cost, local issues, carrier, and site infrastructure plans are in concert with the NGI Goal 2.1; and that support for the NGI applications exists.

The metrics that will be used to determine success of this subgoal are as follows:

1. percent of the 100 sites awarded for connection; and
2. end-to-end bandwidth and service level delivery at the site.

3.1.2.6 Network Management Subgoal

In order to make NGI Goal 2.1 a success, the various testbed operators must be able to ensure that the testbed is manageable and performing to the required levels. To achieve this, network management must be addressed early in the NGI initiative. It is envisioned that the agencies will have distributed, yet coordinated, network management efforts that build on existing agency efforts.

The key features of this network management effort are the following:

1. Distributed help desk
2. Security/authentication methods
3. A distributed gigaPOP Network Operation Center (NOC)

4. Internet 2 and NGI coordination for engineering and network management
5. Capability of applications to schedule end-to-end bandwidth
6. Network monitoring and management tools consisting of existing tools, emerging tools for new services and higher speeds, and scaling of tools to many sites and across carriers

The metrics that will be used to determine the success of this subgoal are as follows:

1. How much control the network operator has of the network
2. The effectiveness of distributed NOCs and help desks
3. Flow measurement and the ability to diagnose problems
4. Development of standards for completely portable performance metrics and diagnostics
5. The effectiveness of the distributed help desk

3.1.2.7 Information Distribution and Training Subgoal

One of the key items necessary to ensure that an initiative succeeds is to provide a mechanism for information dissemination and training. The effort under this subgoal is to ensure that information distribution occurs on a regular basis and that mechanisms are put in place to provide the required level of training to ensure that the NGI can succeed.

These efforts will consist of, but are not limited to the following:

1. Web information dissemination
2. Training opportunities
3. Conferences
4. Agency network/information distribution mechanisms
5. Statistics/measurement information dissemination
6. Interagency program manager meetings

The metrics that will be used to determine success of this subgoal are the following:

1. Number of coordinated conferences held
2. Customer survey (applications and campus network operators)
3. Number of users/customers who indicate (through any of our feedback mechanisms) that they became aware of and engaged with NGI as a result of NGI
4. Information distribution or training activities
5. How many technologies from Goal 1 are implemented into Goal 2.1
6. How quickly applications from Goal 3 use Goal 2.1 networks

3.1.3 Agency Specifics

3.1.3.1 Department of Energy

Introduction. In order to leverage the connectivity of the current ESnet network, the ESnet program will be used to provide connectivity to an additional 10 to 12 DoE-related sites for NGI Goal 2.1. Since DoE has researchers at over 240 universities, the DoE NGI university connections will be accomplished in partnership with the NSF vBNS infrastructure, the NSF Connections program, and the Internet 2 initiative. DoE and NSF have signed a memorandum of agreement delegating to NSF the implementation of DoEís responsibilities for university connectivity using the NSF Connections program and the vBNS and assigning to the DoE responsibilities for intercarrier interconnections described below.

Specifically, DoE proposes to transfer $13 million of the funds DoE receives for Goal 2.1 (university connections) to NSF. NSF will manage the university Goal 2.1 testbed connections program for both agencies in consultation with DoE. DoE will also devote $2 million/year of the funds it receives for Goal 2.1 sites to build and manage next generation interconnections for the Goal 2.1 testbed, in consultation with NSF, the Goal 2.1 testbed sites, and the carriers, to choose and test topology and technology for optimum performance.

Tasks. The DoE task required to accomplish the NGI Goal 2.1 are the following:

1. Create a high performance DoE NGI Virtual Backbone Network (DoE/NGI DN-VBN) over the existing ESnet program commercial infrastructure
2. Interconnect the existing Federal research networks at both the IP and ATM levels via NGI Exchange Points
3. Pursue connecting to advanced NGI-capable gigaPOPs
4. Address QoS issues

A Virtual Backbone. Many NGI applications will be collaborations with universities, industry, and DoE national laboratories. These collaborators and the DoE NGI program will need to use a virtual backbone to meet the high performance requirements of this program without creating a redundant, separate national infrastructure. DoE will use a Virtual Backbone Network to accomplish this goal. The DN-VBN will be structured independently of ESnet, the agency’s mission backbone, but will share the costly underlying commercial ATM communications infrastructure.

A DN-VBN will support the high performance and network research requirements of the NGI program while cost effectively supporting the high performance bandwidth and network research requirements of the applications. Structuring the backbone as a virtual backbone separate from the mission oriented backbone will allow it to operate with routing and policy independent of the requirements of the mission network. Operation of both the virtual backbone and mission network on the same ATM infrastructure allows high performance interconnects between the two without additional latency effects.

The DN-VBN will provide direct connectivity for selected NGI connection members via NGI points-of-attachment which will support IP, ATM, and IP/ATM connections. Other interconnections are discussed below.

NGI Exchange Points: Coordination and cooperation among the NGI Federal member agencies will require high performance interexchange points specific to the NGI program efforts. It is proposed that at least three such facilities be established, one on the East Coast (Washington, D.C., area), one on the West Coast (Ames Research Center), and one in the Midwest (Chicago).

Given that the major activity of the NGI will be based on IP infrastructure with some ATM, it is planned that these three interexchange points will be able to support both IP- and ATM-level interexchange. Currently the major long haul carriers are reluctant to do ATM-level interconnects, and therefore some intermediary is required. In the Chicago NGI Exchange point, NSF and Ameritech Advanced Data Services (AADS) provide the intermediary ATM infrastructure. DoE will provide similar functionality at the D.C. NAP; it is expected that NASA will do so at Ames.

Additionally, since a significant international presence is expected at the D.C. and Chicago exchange points, international collaboration can be supported at those two locations as well. The networking organization Deutsche Forschungsnetz (DFN), responsible for most national research and engineering (R&E) networking in Germany, has already established 90 Mbps of capacity into the D.C. exchange point. Additionally, CERN (European Organization for Nuclear Research) has a connection under way at 4 to 6 Mbps, and several other European entities have expressed interest to connect as well. The NSF international connections program (Science, Technology and Research Transit Access Point–STAR-TAP) is overlaid on the Chicago exchange point, with connections from Japan, Singapore, and Canada already committed or in place, and others from Europe and the Americas in the planning stages.

A block diagram of the D.C. NGI-Exchange Point is shown in figure 4 as one example of architecture for an NGI Exchange Point capable of supporting interconnects both at IP and at ATM levels.

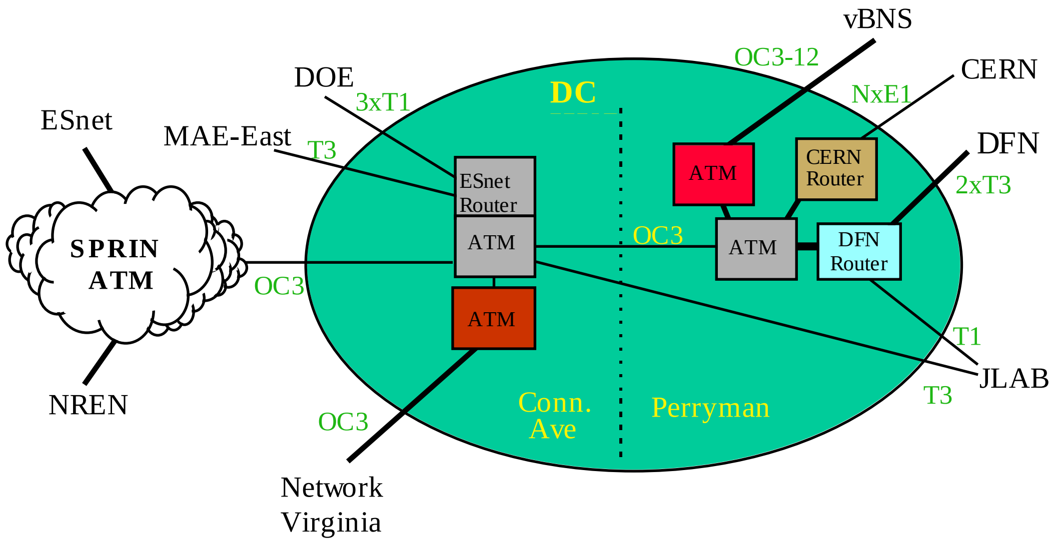

Figure 4. D.C. NGI Exchange Point

GigaPOPs and Interconnection Points. Cooperation among agencies and the high performance focus of the NGI program demand a wide variety of interconnections among participants both in type and in capability. The following connections are planned within the DoE response (see figure 5):

1. *NGI Connection Members.* Connectivity among NGI connection members and DoE principal sites will be through a variety of interconnections depending on performance requirements. High end performance interconnections will use direct connections among the DN-VBN and selected NGI connection members. This may be through direct connections to the NGI connection member associated gigaPOP, or via a direct connection to the member via a DN-VBN NGI point-of-attachment (NGI PoA), or via an NGI Exchange Point. Other connections with less stringent performance requirements will be made by using peering interconnects with other agency NGI networks.
2. *GigaPOPs.* A number of universities will be forming collaboration groups with the equivalent of a regional network established to interconnect the members. Typically, these regional networks will establish one or more peering points for external connectivity, generally called gigaPOPs. DoE may establish physical connections to selected gigaPOPs to meet either mission requirements or NGI requirements. Virtual connections will be established over the physical connection according to ESnet or to the DN-VBN. The university collaboration will allow additional university sites to connect through an NGI-connected university to the NGI. In addition, the universities already connected to the vBNS and ESnet will also be able to connect to the NGI through their existing networks.

3. *Agency NGI Networks.* A number of interconnects with the NGI agency networks are in place, under way, or planned. A number of peering arrangements already exist at the current Network Access Point (NAP) interconnects, including the Metropolitan Area Exchange-West (MAE-West), MAE-East, and the Sprint NAP. A connection to the Chicago NAP is under way.

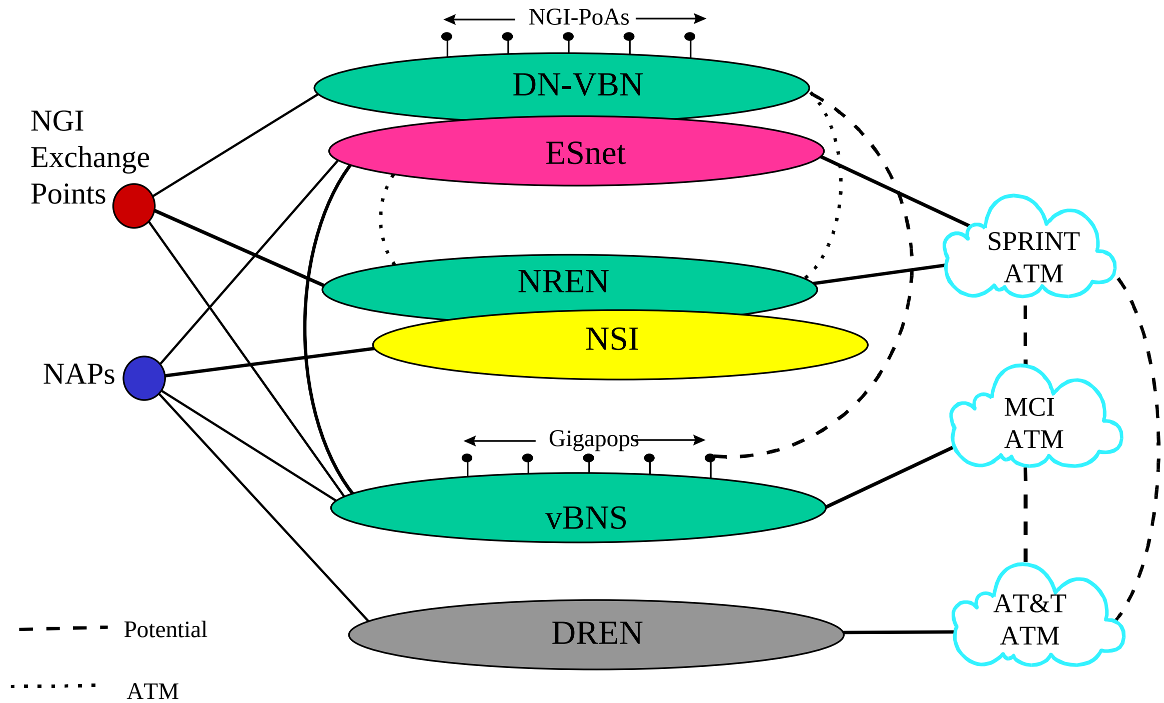

Figure 5: An Initial Architecture for NGI

Other interconnects via the emerging NGI Exchange points are planned. The DC NGI Exchange Point includes interconnects with NREN and vBNS. Other peering arrangements will be established as the other NGI Exchange Points come on line.

Direct interconnects with other agency NGI backbones are also under way. Since the NASA NREN and the DN-VBN reside on the same ATM infrastructure, interconnects between NREN and ESnet or DN-VBN can be established as virtual connections. Such interconnections are already under discussion. Similarly direct interconnects with the vBNS are being implemented. An interconnect in San Diego (between ESnet at General Atomics and the vBNS at San Diego Supercomputing Center–SDSC) was brought online at T3 recently. A second interconnect, at the ATM level, is being planned at the Perryman POP in Maryland, part of the D.C. NGI Exchange Point.

4. *DoE Major Sites* . DoE national laboratories and collaboratory sites are the major components of the membership. ESnet already has a well developed, high performance interconnect structure established for service mission requirements. Incremental upgrades to the ESnet program may be needed to meet the leading edge requirements and traffic demands of the NGI program.

5. *National Carriers.* It is likely that the NGI program will enlist the participation and support of all three major U.S. telecommunications carriers. The NGI will require carrier-to-carrier interconnects among their respective level-2 (data link) communications facilities. Available standards do not yet support such commercial-quality interconnects. This can be managed in the short term by providing intermediate L2 facilities among the carriers. At the D.C. NGI Exchange Point, DoE will provide one such ATM interconnect between the Sprint Connecticut Avenue PoP and the MCI-based Perryman POP.

6. *International.* Although not directly addressed in NGI, international connections already exist in the major Federal networks. NGI interconnections to international collaborations can be made available to NGI-connected sites.

Project Oriented Virtual Networks. Many aspects of the network level research may need to be conducted in a manner that isolates the parameters being investigated. In many cases, a small project oriented network, specific and isolated to the research group, may be established using virtual private connections for the communications among the routers or switches dedicated to the project. This allows the research group the flexibility and independence to investigate various dimensions of the research agenda without affecting other collaborators.

A typical project oriented virtual network would consist of a small number of nodes, of the order of 3 to 12, implemented on behalf of a research project and lasting the life of the project.

Quality of Service (QoS). QoS is an area that will require substantial investigation. The DN-VBN and the project oriented virtual networks will be established, in part, as infrastructure to support the required research and connections, as well as their integration and use by NGI applications. However, initially simple implementations of QoS can be established for the DN-VBN by a variety of methods.

1. *VC/VP Bandwidth.* Virtual circuits and virtual paths can be established with pre-specified bandwidth limits that will give applications a guaranteed level of service. Initially, connections will be set up as permanent connections, with use of switch connections to follow.

2. *Available Bit Rate (ABR) Admission Control.* ABR ports will be established with predefined minimum levels of performance, but with additional bandwidth made available depending on aggregate traffic loads. This will give a guaranteed minimum level of performance, but with higher peak levels available.

3. *ABR Explicit Rate.* This is similar to ABR admission control, but it allows the data that exceed the minimum level of performance to be admitted to the network at an explicit rate that can be accommodated by the network without congestion or cell loss.

4. *Class Based Queuing (CBQ).* CBQ allows for bandwidth allocation at Layer 3 (IPv4) based on a site or type of traffic. When the allocated bandwidth approximations are not in use, the bandwidth may be used by other classes of traffic.

End-to-End Performance. Meeting the NGI program's end-to-end performance goals will require investment in and development of the local infrastructure of NGI member institutions. Although the NGI program is not intended to build permanent infrastructure, it is intended to entice campuses and laboratories to build such an infrastructure so that they may take advantage of the NGI infrastructure.

Funding Issues. The DoE funding strategy for Goal 2.1 will be a combination of funding, service, and equipment grant and loans. Service will be made available in the form of DN-VBN and project oriented networks support, including the general support environment such as interconnections, Domain Name System (DNS), network timing, network management, and help desk. Equipment may be made available in the form of loans or grants to comprise project oriented networks. Advanced interconnection and peering-point management technologies and techniques will be awarded by means of Goal 1 solicitations.

End-Game Transition Strategy. Since the NGI will be built on commercial infrastructure, the transition of this portion of the infrastructure to commercial networks will not be an insurmountable engineering problem. Ultimate success will depend significantly on the business case for the new technologies pioneered on the NGI. Details of the transition strategy on any other portions of the NGI will be worked out at a later date.

Helpdesk. Vital to the success of NGI Goal 2.1 will be an operational Helpdesk, with information and trouble resolution available on a 24-hr daily basis. The scale and scope of the NGI program will necessitate a distributed Helpdesk function that will incorporate support from the centers of responsibility. In particular, trouble resolution must be hierarchical and involve the cooperation of NGI member sites, as well as the sponsoring agency backbone networks.

Milestones and Schedule

Establish DN-VBN, First Phase

FY1998 (1Q)	Establish initial OC-12c User Network Interface (UNI) nodes
FY1998 (2Q)	Roll out routers and switches for the initial phase of DN-VBN
FY1998 (2-3Q)	Upgrade hub-based nodes to OC-12c
FY1998 (4Q)	Start installation of NGI Federal and selected university connections
FY1998 (1-4Q)	Provide initial NGI Helpdesk training; establish trouble ticket system

Install First Project oriented Virtual Networks

FY1998 (2Q)	Demonstrate first project oriented virtual networks on line
Annual	Review proposals, select projects to be supported
On-going	Acquire and install equipment

Install, Enhance Interagency Connections

FY1998 (1-2Q)	Bring West Coast NGI-Exchange into operational status
FY1998 (2Q)	Start online research activities
FY1998 (1-3Q)	Establish connections to other NGI-Exchange points

FY1998 (1-4Q)	Establish connections to international partners
FY1999 (1-4Q)	Upgrade with Early Next Generation Capabilities
FY1999 (1-3Q)	Roll over selected research efforts into backbone QoS
FY1999 (3-4Q)	Demonstrate early end-to-end applications advances
FY2000 (1-3Q)	Upgrade with Final Next Generation Capabilities
FY2000 (1-3Q)	Install final QoS improvement
FY2000 (2-4Q)	Create transition strategy
FY2000 (4Q)	Initiate transition to public sector

3.1.3.2 National Science Foundation

Introduction. NSF projects make a central contribution to NGI goals by leveraging extensive campus and industry partnerships to connect about 100 leading universities with a high performance network fabric, interconnecting this fabric with that of other Federal and foreign research networks, implementing and testing advanced technologies, and supporting hundreds of advanced applications.

The present goals of the NSF programs for high performance connections reflect all aspects of the NGI: high performance connections to about 100 universities and their research partners, with QoS technology, supporting hundreds of meritorious applications.

Tasks. NSF has a two-phase strategy to build fabric for the Goal 2.1 advanced network. The first task will be to significantly expand and enhance NSF's existing program for high performance connections to the vBNS network to serve about 100 leading universities and to link them to their research partners by improving the interconnections among the vBNS and other Federal research networks. NSF will leverage contributions from Internet 2 organization and member universities toward campus and regional infrastructure, as well as national coordination. The second task will be to begin to test and deploy Goal 1 technologies and Goal 3 applications. At the same time, NSF will begin the formation of a formal, national organization of universities to plan and coordinate their role in the NGI and related efforts on an ongoing basis. In Phase 2, NSF will build on lessons learned from Phase I and work with other agencies to design and implement a more unified Federal research network that can better serve the entire research community of interest.

Metrics will focus first on the number and capability of interconnected institutions and later on the extent of successful deployment of Goal 1 technologies and Goal 3 applications.

In Phase I, NSF will connect about 100 leading research universities and their research partners with a high performance fabric by performing the following.

Interconnect the vBNS to Federal research nets

1. Work with DoE, NASA, DARPA, and DoD and other agencies in the Joint Engineering Team to establish optimal interconnection points for the Federal research networks. Build or expand interconnects at SDSC and Ames in California, at Perryman in Maryland, at the Chicago AADS NAP, and at other sites to establish a robust system

that also supports efficient routing. Focus on the NSF supported STAR-TAP in Chicago as an interconnection point where Federal networks can exchange research traffic with similar networks from other countries.

2. Support high performance gigaPOPs and other interconnects at about 20 locations designated by connected universities and the Internet 2 organization in addition to the Federal interconnects. Coordinate Internet 2 and Federal interconnects for interoperability and shared experimentation with NGI technologies and applications.

Interconnect the fabric to foreign high performance research nets

1. Model future interconnections on the NSF agreement with Canada's Research and Engineering Network CA*net II that identifies specific institutions participating in the research partnership for high performance networking. (Test the use of Border Gateway Protocol (BGP) communities, tag switching, or other technologies to identify and route the high performance traffic of the designated sites.)

2. Use the new NSF High performance International Internet Services grant program (NSF 97-106) to leverage the interests of other nations. Request consortium proposals for partial NSF funding from foreign networks and institutions working with U.S. institutions. Suggest interconnection for research partnerships at the STAR-TAP, that is, the "Science, Technology and Research Transit Access Point" (http://www.startap.net), at the Ameritech Advanced Data Services Network Access Point in Chicago.

3. Expand the NSF supported STAR-TAP to support additional high performance interconnections of Federal and foreign research networks. Use this "common point of contact" approach to resolve problems of support for transit traffic and of issues surrounding multiple AUPs (acceptable use policies).

Complete campus connections to the vBNS

1. Five supercomputer centers and 59 campuses were awarded partial funding for connections in FY1997. Each raised at least 50 percent matching funds locally, and pledged to support the resulting high performance connections for an indefinite period after the 2-year funding of the award. Most are connecting to the NSF vBNS national backbone at DS3 and OC-3 rates.

2. Award about 35 new connections in FY1998, using NGI funding to increase awards to achieve the more aggressive performance goals of the NGI. NGI funds will enable the connection of the targeted institutions at higher speeds and earlier than would be possible with existing NSF funding alone. This improvement is a key requirement for the overall NGI program Goals 1 through 3.

3. Upgrade the vBNS fabric as often as is feasible (now operating at OC-12). The vBNS cooperative agreement fits NGI Goal 2.1 very well. It calls for a "leading edge but stable" network that is always at a level of performance beyond what can be purchased on the market. It will be upgraded, accordingly, as newer technologies become available, and must be upgraded initially to support the larger number of institutions and interconnects called for in the NGI.

4. Upgrade existing campus connections as required to meet NGI goals. Institutions now connecting at DS3 rates must be upgraded to OC-3 to meet the NGI performance objectives. This will require additional funding of NGI, in some cases, and the continued development of regional and campus infrastructure in many locations.

Supporting access from rural or remote sites

1. Use the EPSCoR program (Experimental Program to Stimulate Competitive Research) to increment NSF awards for university connections in certain states in which network access is more difficult or expensive. (The EPSCoR states are: Alabama, Arkansas, Idaho, Kansas, Kentucky, Louisiana, Maine, Mississippi, Montana, Nebraska, Nevada, North Dakota, Oklahoma, South Carolina, South Dakota, Vermont, West Virginia, and Wyoming, and the Commonwealth of Puerto Rico.)
2. This program can provide an additional $200,000 supplement to EPSCoR institutions that have otherwise met all the criteria for a high performance connections.

Work with the Internet 2 organization and others to form a national organization of leading universities to help coordinate the future development of the advanced networking fabric:

1. NSF supported workshops over the past 3 years have led to the definition of university requirements for advanced networking and have facilitated the organization of over 100 leading universities as “Internet 2.”
2. Internet 2 now represents essentially the same constituency as that of the NSF high performance connections program.
3. NSF will work with Internet 2 and other stakeholders to develop a new, formal organization representing universities committed to the active research and development of advanced networks. This organization, which might be modeled on similar organizations that operate other national facilities for research such as UCAR (the University Corporation for Atmospheric Research; http://www.ucar.edu/), will represent the research and development needs of the universities involved. The new organization will play an essential role in the national coordination of advanced networking projects and activities such as the following.

Design and manage compatible gigaPOPs to test and support network interoperability

Coordinate national scale development and testing of advanced network technologies for QoS, security, measurements, etc.

Facilitate national scale applications among the universities and their research partners

In Phase 2, NSF will work with partner universities and with other Federal agencies to implement a more unified, high performance network fabric:

1. Build on lessons of previous tasks
2. Coordinate system of gigaPOPs and other interconnects
3. Implement managed, interoperable QoS and other services across connection fabric

Milestones.

FY1998 (1Q)	Interconnect the vBNS with Federal research nets
FY1998 (1Q)	Facilitate initial national scale applications
FY1998 (2Q)	Complete about 100 campus connections to the vBNS (building on 29 in 2Q 1997, 35 more 4Q 1997, 35 more apply 3Q 1997)
FY1998 (2Q)	Form a national organization of leading universities to coordinate the development of the QoS fabric
FY1998 (3Q)	Interconnect the vBNS to foreign research nets
FY1998 (4Q)	Coordinate national scale testing of advanced network technologies
FY1999 (1Q)	Coordinate compatible gigaPOPs for QoS, etc.
FY1999 (2Q)	Upgrade remaining campus connections to OC-3 and above
FY1999 (2Q)	Facilitate national scale QoS applications
FY1999 (4Q)	Implement unified network structure with partner universities and agencies

3.1.3.3 National Aeronautics and Space Administration

Strategy. NASA will leverage the NASA Research and Education Network (NREN) in meeting its NGI goals. NASA will provide both a high performance network application testbed and a network research testbed for the NASA community and its partners. These testbeds exist at the various NASA centers now and can be interconnected via NREN thus providing virtual testbeds and harnessing the expertise distributed throughout NASA. In NGI Goal 2.1, NASA will focus on delivering a leading edge application environment. Therefore, NASA will (1) enable next generation application demonstrations across the network; internetwork with other Federal agencies and academic and industry partners at both the IP and ATM service level; and deploy advanced networking services such as IPv6, multicast, QoS, security and network management tools.

Enabling Applications. The NASA community and its Federal partners have many applications that will require the facilities of the NGI to be fully successful. Specifically, they will require access to a high performance network that is compatible with the current Internet. NASA plans to leverage the NREN ATM infrastructure to provide a high performance network to NGI application partners.

Internetworking (NGI Exchange–NGIX– and GigaPOPs). Coordination among the Federal networks and the university-initiated gigaPOPs is a crucial element to NGI success. This will involve developing and implementing an internetworking architecture among the network and their WAN service providers, namely, Sprint, MCI, and AT&T.

Three facilities for interexchange will be established initially to support internetworking: one on the East Coast (D.C.), one on the West Coast (Silicon Valley, Ames Research Center) and one in the Midwest (Chicago). These exchanges will support both IP and ATM bearer services. NASA will lead the development and support of the West Coast Exchange, while DoE and NSF will lead the development and support of the East Coast and Midwest exchanges, respectively.

Additionally, NASA will peer with universities through the proposed Internet 2 gigaPOP architecture. Initial interconnections are planned at the Midwest Exchange and the West Coast Exchange with others to be added as application requirements dictate.

NASA will leverage its experience in high speed satellite data communications from the Advanced Communications Technology Satellite (ACTS) program and attempt to make use of existing NASA satellite resources, as well as seek out satellite services from commercial sources. These high speed links could provide a means of connecting international testbeds to the NGI (e.g., GIBN— Global Interoperability Broadband Network) (see fig. 6).

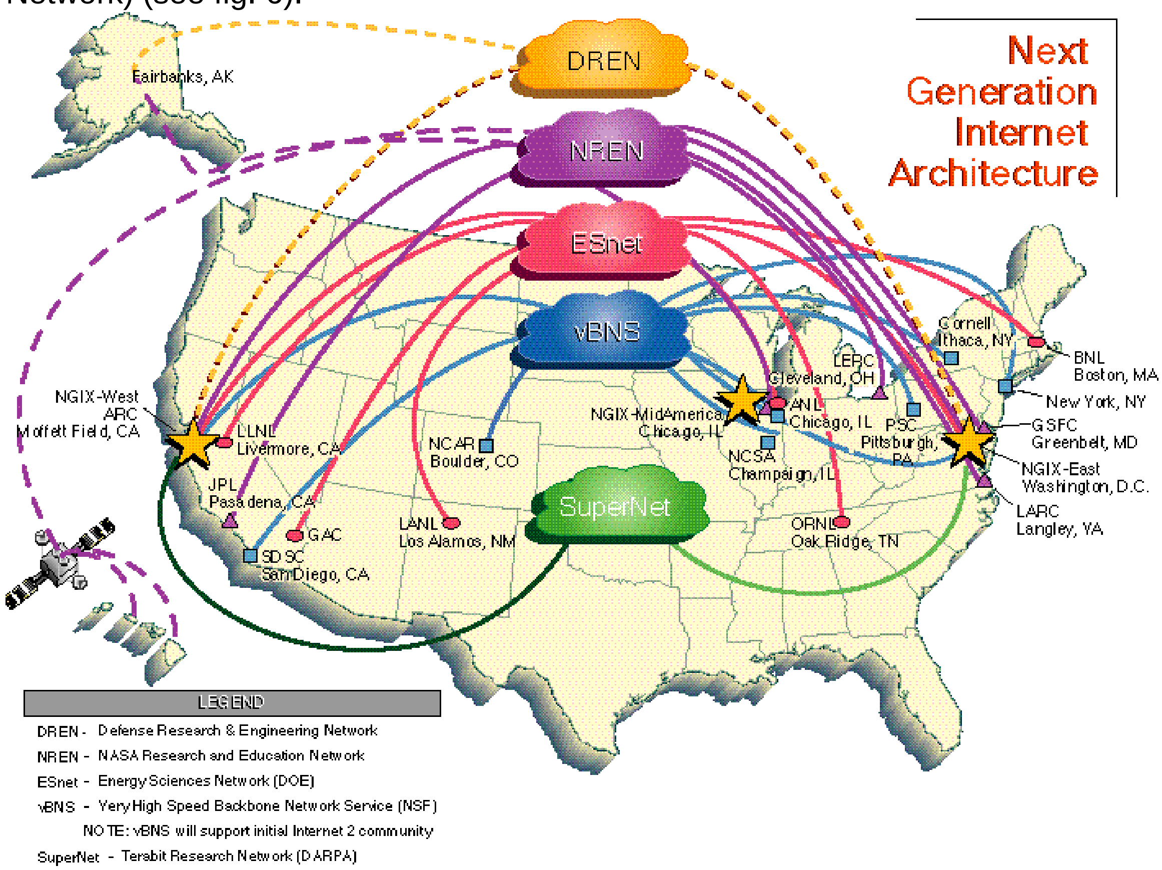

Figure 6: Proposed NGI Architecture.

Advanced Network Services. Several advanced network services will require extensive research as they are carried out under NGI Goal 1. Goal 2.1 will rapidly adopt these technologies as they are proven in order to bring them to the NGI application community. These advanced network services can be validated in the virtual network research testbeds. Interoperability testing will be paramount to roll out these new services.

Initially, NASA's network will be IPv4, native multicast, best effort network over ATM to the end user application community. The network will rapidly evolve to offer additional addressing, priority, and management services.

Milestones and Schedule:

Enabling Applications

FY1998 (3Q)	Implement experimental OC-12 service at three sites
FY1998(4Q)	Upgrade selected NGI users to end-to-end 10+ Mbps LAN access
FY1998 (4Q)	Ensure minimum of three NASA sites at production OC-12
FY1998 (4Q)	Test and implement Available Bit Rate Service
FY1999 (1Q)	Test and implement Constant Bit Rate Service
FY1999 (2Q)_	Test and implement Switched Virtual Circuits for bandwidth on demand
FY1998 (1Q) FY1999 (4Q)	Establish international high performance connections supporting application partnerships
FY2000 (2Q)	Implement 100+ Mbps LAN access to three NREN users
FY1998 (1Q) FY2000 (4Q)	Support testing for efficient, low bit error rate interfaces among terrestrial nodes, satellites, and mobile wireless networks
FY1998 (1Q) FY2000 (1Q)	Support the development and validation of hybrid communications architecture models; leverage activities performed under NREN and program
FY2000 (4Q)	Demonstrate with 100+ Mbps end-to-end communications over wireless and wireline networks
FY2002 (4Q)	Provide at least 100+ Mbps to end users' desktops in a wide area environment by providing OC-48 (2.5 Gbps) service to support collaborative multimedia applications

Internetworking (NGIX and GigaPoPs)

FY1998 (3Q)	Interconnect NASA testbeds with at least two NGI Partners
FY1998 (3Q)	Scope and design network of networks architecture and network management and control with Federal partners; leverage partner investments to provide target OC-3 connection to sites using an interagency OC-12 backbone

FY1998 (4Q)	Upgrade NGIX to OC-12 capability
FY1999 (1Q)	Interconnect NASA to two other NGI networks at OC-12
FY1999 (4Q)	Test network-to-network links at OC-12
FY1999 (4Q)	Establish will-carry and peering arrangements with Federal research networking partners
FY1999 (4Q)	Establish peering arrangements at NGI Exchange Points
FY1999 (4Q)	Establish peering arrangements at GigaPOPs
FY1999 (4Q)	Establish cross-agencies collaboration strategy and cost sharing agreements
FY1998 (1Q) FY1999 (4Q)	Interconnect at least one NASA site and at least one university facility to ACTS
FY 1998 (1Q) FY2000 (4Q)	Cooperate with international networks as appropriate to meet the needs of NASA partners and address NASA international connectivity requirements
FY1998 (2Q) FY2000 (4Q)	Interconnect international WANs to NGI Exchange Points
FY2000 (4Q)	Internetwork NASA testbeds with other Federal agency networks to create an interoperable interagency network of networks
FY1998 (1Q) FY2002 (4Q)	Pursue private sector satellite service partners to further technical achievements of ACTS program (esp. High Data Rate program— HDR).

Advanced Network Services

FY1998 (3Q)	Implement next generation network management and monitoring for the NASA testbed
FY1998 (4Q)	Interconnect NASA networks with select broadband links to identify and evaluate network management and control, security, interoperability, and other technology issues
FY1998 (4Q)	Interconnect NASA testbed to five NASA scientific and research LANs
FY1998 (4Q)	Implement secure network technology across NASA testbed(s)
FY1998 (1Q) FY2000 (4Q)	Build a virtual NASA testbed through collaborative efforts of existing NASA centers
FY2000 (2Q)	Implement Layer-2 security technology at five NASA sites
FY2000 (4Q)	Implement IPv6 security technology at five NASA sites
FY2001 (4Q)	Implement network security technologies and policies across NASA testbeds
FY2001 (4Q)	Implement IPv6 security technology on NGI Exchanges

FY2002 (4Q)	Develop and implement interagency security policies

3.1.3.4 Department of Defense

The following paragraphs describe participation of the DoD High Performance Computing Modernization Program (HPCMP) and its constituent Defense Research and Engineering Network (DREN) in the NGI. The degree of participation will be commensurate with NGI funding arrangements: The Level I cooperation proposal assumes no NGI funds; the proposed Level 2 and 3 activities, participation and collaboration, will require NGI funds for any new NGI-specific sites that are established. Distinctions between Level 2 and Level 3 efforts include (1) Level 2 contributes the DREN Intersite Service Contract (DREN/ DISC) production fabric as a transit medium of convenience, and (2) Level 3 expands the scope of the Advanced Communications Technology Satellite (ACTS) ATM Internetwork to include NGI collaborators.

Level I: Cooperation with NGI Collaborators. Establish gateway connectivity among DREN and select aggregation points, such as the STAR-TAP node in Chicago through the DREN/DISC. Support access among NGI-connected sites and DREN sites is through this gateway connection. DREN sites are typically DoD laboratories, DoD-sponsored high performance computing centers, and affiliate universities. DISC may provide additional gateway connections at locations of natural affinity, such as DoD HPCMP's Programming Environment and Training (PET) affiliates. To maximally leverage NGI outreach, particularly to the Internet 2 community, these connections may be attached to the nearest gigaPOP.

Level 2: Participation Through Provisioning Selected New NGI Nodes. With NGI funding, HPCMP will support IP or ATM delivery services to university or agency sites through the DREN/DISC contract. It is anticipated that these will typically be sites near DREN nodes, thereby benefiting from shorter access distances and reduced access carrier costs. To meet Goal 2.1 objectives, these would initially have OC-3 services. For NGI client sites, DREN/DISC would provide service delivery points and transit carrier fabric for subscriber network traffic to reach other service delivery points, as well as possible further routing across collaborator networks to other target destinations.

Level 3: Collaboration via the DREN Testbed/ACTS ATM Internetwork (AAI). HPCMP will support subscription of new NGI-sponsored sites to the ACTS ATM Internetwork (AAI). These NGI sites will be encouraged to collaborate with HPCMP participant sites or AAI research partner sites (i.e., DARPA-funded nodes). Principal objectives of this arrangement include advancing the AAI high performance network research agenda, as well as supporting applications or experiments requiring high performance network resources.

Principal near-term network research objectives include establishing Peer Network-to-Network Interface (P-NNI) hierarchy across ATM domains, network performance measurement, congestion management, IP and ATM address resolution mechanisms, and ATM signaling behavior across multiple providers. Another way to add new AAI nodes is to establish gateway agreements with other providers, such as through NSF for selected vBNS attached collaborator organizations.

Milestones:

FY1998-2001 Establish and support gateway connection to STAR-TAP
FY1998-2001 Establish additional gateway connections as appropriate
FY1998-2001 Add selected sites to DREN/DISC
FY1998-2001 Add selected collaboratory sites to ACTS ATM Internetwork

3.2 GOAL 2.2: NEXT GENERATION NETWORK TECHNOLOGIES AND ULTRAHIGH PERFORMANCE CONNECTIVITY

3.2.1 Introduction

This Goal 2.2 addresses the development of ultrahigh speed switching and transmission technologies, and the demonstration of end-to-end network connectivity at 1+ Gbps. Because of its high risk and pioneering nature, networks involving about 10 NGI sites and applications will be implemented. Attaining this goal, together with the technologies developed in Goal 1, will be the pathway to terabit-per-second (Tbps) networks, operated with the appropriate network management and control and guaranteed end-to-end QoS. Working in partnership with industry is the key to a shared infrastructure that can be profitably used to support high end scientific users and large numbers of commercial users.

This is a joint agency effort with DARPA as the lead and with participation by DoE, NASA, NSF, and other Federal agencies.

3.2.2 Strategy

The technology of choice to achieve a Tbps network is wavelength division multiplexing (WDM), which is a technique of mixing many wavelengths onto the same optical fiber. This is equivalent to opening up the narrow communication links into multiple-lane communication highways. DARPA’s Broadband Information Technology (BIT) program has pioneered much of today’s WDM effort. Whereas WDM is currently implemented at the physical layer, the aim here is to integrate WDM and its management with the upper layers of ATM and IP. To this end, a new network architecture that specifically addresses the access loop will be designed.

These technologies and architectures hold the promise of eventually satisfying the goal of an infrastructure that is shared by both high end users, typical users, and network researchers. Some of the network nodes will be chosen to coincide with some Goal 2.1 nodes. The architecture will be designed such that portions of the Goal 2.1 networks can interconnect to and gracefully evolve into Goal 2.2 networks and further demonstrate the continuing evolution of network performance. Partnership with long distance and local exchange carriers is a key to ensuring the early adoption of this technology and to ensuring its affordability.

With respect to the generation-after-next network technologies, Goal 2.2 will explore optical, electronic, and hybrid switching techniques. The goal is to pave the way to Tbps packet switching systems. On the optical side, hybrid dense WDM and optical time division-multiplexing (TDM) systems will be explored. On the electronic side, a distributed electronic switching design, as opposed to a single monolithic Tbps module, will be pursued. Resulting devices and systems will be initially field tested in the Goal 2.2 network. These research activities will be a combination of government, academia, and equipment vendor efforts and collaboration.

3.2.3 Metrics

Goal 2.2 will focus first on the deployment of at least one metropolitan network (e.g., five-node network), with the appropriate management and control software. This network will be operational at least 80 percent of the time and will be capable of delivering 20 Gbps to each node. As the tools of Goal 1 and the broadband local access technologies become available, they will be incorporated into this network to experiment with providing ultra-high speed end-to-end QoS, management of lead user infrastructure, data integration, and network security. A true test of the success of a network is the range of new applications it will enable. To this end, at least 10 new applications will be tested on this ultrahigh speed network.

In Phase 2 of Goal 2.2 the network will be expanded to a wide area network with about 10 nodes performing similar functions, as in Phase I. In this phase, agency applications will be linked to demonstrate a distributed, heterogeneous, multidomain, and multivendor environment. Since the number of nodes that can be built is limited by the available resources, the scalability and network management of hundreds of ultrahigh speed nodes will be examined by simulation and modeling.

The following is a description of the implementation plans and milestones by participating agencies.

3.2.4 Agency Specifics

3.2.4.1 Defense Advanced Research Projects Agency

Wide Area Broadband Core: DARPA's Broadband Information Technology (BIT) program has developed basic WDM transmission capabilities and will soon demonstrate a metropolitan network of five nodes, with link transmission capacities of 20 Gbps. DARPA will extend these technologies and deploy them in more complex, mesh-like topologies that involve long distance links.

The metropolitan testbed will be expanded into a wide area network with about 10 nodes using WDM technology. This wide area backbone will have sufficient aggregate transmission and switching resources to support hundreds of users at Gbps rates. This network will share the fiber facilities with the general public.

Tbps Multiplexing and Switching: DARPA will develop the generation-after-next multiplexing, switching, and routing technologies that will bridge the gap among packet-based Gbps tributaries and the WDM-based optical core. This task will also lay the groundwork for the direct optical support of packet-based communication. A major component of this task will be to investigate statistically sound techniques for performing "space-division"-like spreading of the resultant TDM traffic across a set of wavelengths. A second component will be the design and demonstration of a highly parallel and distributed switching fabric.

Taken together, these efforts will enable the development of a highly distributed approach to Tbps switching, based on a combination of optical and electronic technologies, with many-to-many multicast capability.

Broadband Local Trunking: The need to provide selected sites with "orders-of-magnitude-above-average" access to the network core has been a recurring source of delay in commissioning advanced research facilities. This task will explore novel and cost effective approaches to delivering broadband access to selected sites within a geographically restricted area. DARPA will examine the terrestrial extension of SuperNet rate facilities to the building and explore the effectiveness of high capacity (>150 Mbps) radio frequency (RF)-based trunking. In addition, wireless broadband local access will be addressed as one of the DARPA tasks in Goal 2.2.

Technology Demonstration and Field Trials: Most of the technologies to be developed by the previous tasks are associated with the physical, link, and networking layers. This task will seek opportunities to demonstrate the newly developed capabilities through collaboration with some of DARPA's application oriented activities, such as the Human Computer Interaction, Information Management, and Intelligent Collaboration and Visualization programs.

Milestones:

Wide Area Broadband Core

FY1998	Simulate WDM transmission in WAN
FY1999	Demonstrate five nodes WDM WAN at 2.5 Gbps per channel
FY2000	Demonstrate five nodes WDM WAN with arbitrary add/drop channels
FY2001	Establish five nodes WDM MAN at 10 Gbps per channel
FY2002	Establish 10 nodes WDM WAN with 160 Gbps facilities

Tbps Multiplexing and Switching

FY1999	Install 300 Gbps electronic ATM switch
FY2000	Install hybrid ATM/WDM burst switch
FY2002	Test Tbps packet switching

Broadband Local Trunking

FY1999	Test WDM broadband local access architecture
FY1999	Test broadband wireless trunking and networking
FY1999	Test WDM local access network elements
FY2000	Demonstrate WDM access network at 1 Gbps
FY2002	Demonstrate 1 Gbps end-to-end access in WAN

3.2.4.2 Department of Energy

DoE will ensure that at least two DoE sites are active participants in the extremely high speed gigabit (and up) NGI Goal 2.2 testbeds led by DARPA, so as to enable DoE and interagency advanced applications and users. This will enable the investigation of signaling, data access service guarantee mechanisms, protocols, gigabit speed congestion and flow control mechanisms, especially efficient and capable operating system support for high end workstations, parallel systems, and storage servers. DoE will also continue its leadership in extremely high speed end-system interfaces and protocols, as well as their integration with the operating systems, since these are crucial to the success of DoE's mission and interagency advanced applications. DoE will accomplish the following:

1. Connect at least two DoE sites to the DARPA led 1000x testbed(s) to investigate gigabit switching and congestion control for IP and other protocols (e.g.,

ATM or point-to-point protocol—PPP) as appropriate, and to demonstrate end-to-end gigabit applications

2. Evolve some of the high end DoE applications to use these technologies and capabilities
3. Develop analyzers, premise switches, and end system interfaces for advanced 1000x-class end system (e.g., HIPPI64) interfaces and protocols
4. Develop and deploy 1000x-capable network management and monitoring tools
5. Provide interconnection among the 1000x to 100x networks or sites
6. Develop and deploy aggregation/deaggregation network capabilities for lower speed tributaries (e.g., OC-12, OC-3, DS-3, etc.) so as to provide for testing and analysis of the advanced higher end services (gigabits) and technologies, as well as provide a transition path to these new services

Milestones:
DoE has identified the following possible areas to support Goal 2.2, but will also pursue other areas that have not yet been identified, as well as leverage DARPA and other agency programs.

FY1998	Develop HIPPI64 analyzer
FY1999	Develop premise switch
FY1999-2001	Provide API for striping control and data
FY1999-2002	Provide interface(s) for high end parallel systems
FY2000-2003	Provide interface for SONET and WDM

DoE will develop aggregation/deaggregation techniques for OC-3 into OC-12 and OC-12s into gigabit capabilities for testing and development purposes:

FY1998-2000	Develop techniques for SONET
FY2000	Connect at least two DoE laboratories to a 1000x testbed
FY1999-2002	Develop techniques for WDM

3.2.4.3 National Science Foundation

NSF will participate actively in NGI Goal 2.2 through selected connections to the ultrahigh speed networks as well, as the direct funding of competitive research proposals by campus-based investigators. NSF will participate with DARPA and other agencies in ultrahigh speed networking links and technologies through NSF's two major supercomputer partnerships (Partnerships for Advanced Computational Infrastructure, PACI) centered at San Diego Supercomputing Center (SDSC) and the National Center for Supercomputing Applications (NCSA). The focus will be on protocols and technologies for advanced, distributed computing. NSF strategies will center first on peer evaluation of Goal 2.2 research and new end-to-end network technologies, and later on the deployment of Goal 1 technologies to Goal 2.2 networks. Among other activities, NSF will do the following:

1. Connect two PACI supercomputer sites to Goal 2.2 networks

2. Select and tune PACI applications for high speed research
3. Study and tune ultrahigh speed performance using future generations of tools such as OC12-MON
4. Connect to applications at selected partner universities
5. Provide selective interconnection among the 1000x to 100x networks
6. Adapt Goal 1 results to Goal 2.2 networks
7. Coordinate these activities through the NSF National Laboratory for Applied Networking Research and the PACI program, as well as awards to individual PIs

Milestones

FY1998	Select and tune PACI applications for high speed research
FY1998	Award additional peer-reviewed, campus-based projects
FY1999	Connect two supercomputer sites to Goal 2.2 network
FY1999	Study and tune high speed performance
FY1999	Connect to applications at selected partner universities
FY2000	Integrate results of campus-based research
FY2000	Adapt Goal 1 results to Goal 2.2 networks

3.2.4.4 National Aeronautics and Space Administration

Strategy*:* NASA will partner with DARPA to have at least two NASA sites be active participants in ultrahigh speed testbeds. NASA will investigate the feasibility and performance of engineering application demonstrations across these testbeds. The goal is to achieve an end-to-end high speed hybrid network capable of supporting both wireless and bounded media applications.

Milestones*:* The goal of NASA's program in collaboration with NGI is to accelerate R&D in selected core technologies (transmission, fast switching, wavelength division multiplexing, network security). The following milestones illustrate NASA's participation:

FY1998(3Q)	Establish a system of high performance network interconnection points in partnership with industry and academia; provide for vendor neutral connection and access to high performance networks; provide direct access to very high speed experimental applications and facilities; interconnect to at least two gigaPOPs and one industrial testbed
FY1998(4Q)	Partner with sites that are experimenting with multigigabit networks in selected laboratories, campuses, and regions to establish high speed networking research testbeds
FY1998(4Q)	Negotiate collaborative agreements with at least five industry partners
FY1998(4Q)	Establish R&D plans with industry partners on advanced network technologies in switching and routing

FY1999(3Q)	Partner with industry to test and develop aggregation/deaggregation techniques for OC-48 service
FY1999(4Q)	Demonstrate optical and fast switching networks; connect two NASA sites to the DARPA Broadband Information Technology testbeds
FY1999(4Q)	Test network-to-network links at OC-48
FY1999(4Q)	Scope and develop with industry the overall system cost, cost share, and collaboration to enable transfer of appropriate technologies
FY2000(3Q)	Implement experimental OC-48 service at three sites
FY2001(1Q)	Develop application performance benchmarks for gigabit and terabit testbeds.
FY2002(1Q)	In partnership with industry and academia, develop performance measurements for OC-48
FY2002(4Q)	Connect two NASA sites at OC-48 and adapt applications and report performance
FY2002(4Q)	Scope and design network of networks architecture and network management and control with Federal partners; leverage partner investments to provide target OC-12/48 connection to sites using OC-48/192 as an interagency backbone; develop aggregation/deaggregation traffic schemes.

4. GOAL 3: REVOLUTIONARY APPLICATIONS

4.1 INTRODUCTION

Applications are the ultimate success metric of this program. Faster and more advanced networks will enable a new generation of applications that include crisis response, distance education, environmental monitoring, health care, scientific research, and national security.

To achieve this goal, agencies will leverage the NGI investments significantly with other major application investments. Agencies will demonstrate new applications, as well as enhance and enable current mission applications that meet important national goals and missions. Each demonstration will partner the advanced networking technologies developed in Goals 1 and 2 with modern applications technologies. Each community will bring its knowledge, skills, and methods to the partnership. The applications partner will provide the bulk of the resources and support needed to implement its applications. The partner will work within the framework of the NGI initiative to develop and demonstrate its applications over the high performance networking infrastructure by using advanced network technologies provided by other parts of the NGI. The applications demonstrations will primarily be proof-of-concept demonstrations. Many of these demonstrations will suggest new ways for the application partners to meet their mission needs. The demonstrations are part of a research effort:, as such, they will initially be built on less-than-fully-robust technologies and be operating in less-than-bulletproof networking environments.

Many agencies have critical signature applications that will benefit from advanced networking services and capabilities. Both the Federal governmentís information technology services and the Federally supported R&D community have networking requirements that cannot be met with todayís networking technology. Higher speed networks with more advanced services and functionality will enable a new generation of applications that support fundamental governmental interests including disaster response; distance education; environmental monitoring, prediction, and warnings; national security; scientific research; and health care.

As the NGI initiative develops capabilities such as QoS, nomadicity, and adaptive networking for the NGI, advanced demonstration applications will take advantage of these new services. It is expected that additional agencies will participate in these applications. For example, the National Institute of Health (NIH) has already indicated its intention to join the initiative, and the Environmental Protection Agency (EPA) and the National Oceanic and Atmospheric Administration (NOAA) have identified key applications requiring NGI speed and services. The education community is putting significant effort into connecting K-12 schools to the current generation Internet. Advanced education applications such as distance learning are expected to be major partners as the NGI matures. NGI applications prototypes will test these new capabilities to ensure that the protocols developed in Goals 1 and 2 are complete, robust, and useful in real applications and to provide a road map to future governmental and commercial services.

Success in reaching Goal 3 depends on success in Goals 1 and 2. Hence, Goal 3 drives the selection of the capabilities and designs for the other goals. In addition it requires integrating the new networking capabilities with the application domain.

Although this program will not provide substantial direct funding for applications, it will partner with and leverage resources of the NGI community to incorporate new networking technologies and capabilities into applications to improve the R&D and service delivery to the public and private sectors. The essential and common features required by applications and demonstrated by this program will be identified and included in the feature set available via the NGI.

4.2 APPLICATION SELECTION AND COORDINATION

4.2.1 NGI Applications Selection Process

NGI applications are selected through four interrelated applications identification processes, each with its own rationale, selection criteria, and funding approaches: (1) NGI funded agency missions, (2) NGI affinity groups, (3) Committee on Computing Information and Communications (CCIC) Applications Council, and (4) broader communities. These four processes together result in lists of candidate applications that are then sorted and ultimately selected according to criteria including resource requirements and benefit to the NGI program, the funding agencies, and the Nation. The overall NGI applications selection process is described in the following sections (see fig. 7).

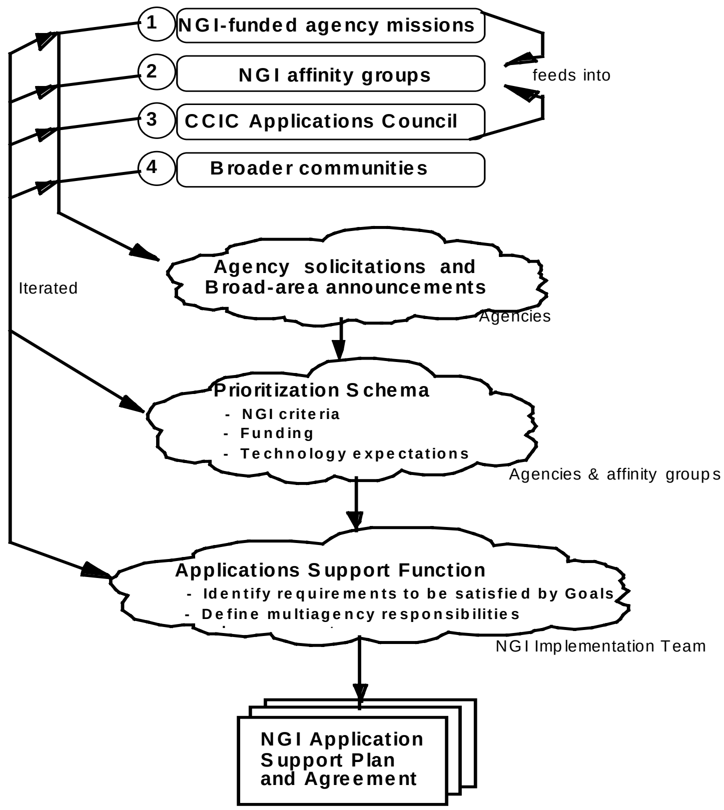

Figure 7. NGI Application Selection Process.

4.2.1.1 NGI funded Agency Missions

Each participating agency identifies its mission applications of interest. These applications are important to the agency mission and are such that the NGI's increased networking functionality and performance are needed in order for the applications to be implemented. These application candidates are identified within the agencies and brought to the NGI program by agency champions. NGI funded agencies are also responsible for outreach to other CCIC R&D agencies that are outside the initial set of NGI funded agencies, to encourage these other agencies to identify, fund, develop, and demonstrate their own mission applications within the context of NGI. A key means for this outreach is agency participation in NGI discipline affinity groups.

4.2.1.2 NGI Affinity Groups

NGI applications are coordinated by NGI affinity groups. These groups are identified in discipline areas such as health care and environment, and also in technology areas such as collaboration and remote operations. Each group will coordinate applications development for its area of interest, as well as harmonize the requirements its applications need, from NGI Goals 1 and 2.

The affinity groups are responsible for outreach to their own communities of interest; for example, the health care affinity group would not focus solely on coordinating the health care related applications brought by its agency participants, but would spend some effort to identify the best candidates throughout its discipline area and to reach out to the broader health care community in search of the best health care related applications for the NGI. The result is the identification of applications that are proposed for formal NGI selection. More detail on the NGI affinity groups is provided in subsequent sections.

4.2.1.3 CCIC Applications Council

The CCIC Applications Council reports to the White House National Science and Technology Council (NSTC) Committee on Computing, Information, and Communications (CCIC) which seeks to provide outreach to the non-CCIC R&D agencies. The Applications Council will help to identify applications that are input either to NGI affinity groups or to mission agencies for coordination and support. Affinity groups may also be established at the CCIC Applications Council's suggestion.

4.2.1.4 Broader Communities

The processes are intended to serve the interests of the Federal agencies in carrying out their legislated mandates. However, it is recognized that many innovative NGI applications will come out of non-Federal communities, or at least will not be identified by the applications identification processes identified above. For example, the World Wide Web was developed initially with government funding but did not derive specifically from mission requirements, and it had an effect far beyond the expectations of the funding agencies. Many other successful applications are developed completely without government funding.

Several opportunities exist for the best applications, independent of source, to be proposed as candidates for NGI selection, including solicitations by agencies using procedures such as Broad Area Announcements (BAAs) and Cooperative Agreement Notices (CANs), as well as outreach programs coordinated with broader communities such as Internet 2 and Highway 1.

4.2.2 Funding

Applications are funded by a variety of means, consistent with the four applications identification processes above. It is understood that most of the resources applications must come from sources outside of the NGI program itself. The contribution by NGI to these applications is intended to enable the applications to be implemented within the NGI framework and demonstrated within the NGI context. Nevertheless, NGI will play an active role in identifying applications requiring NGI resources. NGI funded mission agencies will have specifically targeted funding that is earmarked for supporting their advanced program needs.

Applications coordinated by NGI affinity groups will be worked within the participating agencies and organizations to identify the funding approach that will provide adequate resources for the development of the application. Broad community announcements will specify the amount of funding or cost sharing available for applications development. NGI funding will be available from the NGI funded agencies to meet most if not all of the NGI support requirements of the applications, including network connectivity provided within the resource limitations of Goal 2 and including network services and technologies developed by Goal 1. The agreement of Goals 1 and 2 to provide the needed support is part of the final selection and go-ahead process (see Applications Support Function below).

4.2.3 Prioritization Schemata

Applications proposed as candidates by any of the four applications identification processes described above will be evaluated for funding and implementation mainly by three areas of characterization: NGI criteria, funding approach, and technology expectation.

1. *NGI criteria* represent the extent to which the application fulfills the identified purposes of the NGI program, namely., to showcase applications which meet important national needs and which require the success of Goals 1 and 2 in order to be implemented. The NGI criteria are elaborated below.
2. *Funding approach* represents the extent to which adequate funding will be assured, both funding for the application development itself, as well as funding for the integration of the application within the NGI network, demonstration of the application, and carrying out of the technology transfer approach if the application demonstration is successful.
3. *Technology expectation* is the degree to which the application-defined technology plan satisfies the NGI technology requirements, including the defined approaches for demonstration, evaluation, validation, scalability, and deployment or commercialization of the resulting capability.

Agencies, the engineering teams of Goals 1 and 2, and the affinity groups will apply these prioritization schemata iteratively to help define and clarify the specifications and ultimately make the final selections of the applications for NGI implementation.

4.2.4 NGI Criteria

This program will select partnerships for application demonstrations and testbeds which not only meet critical governmental needs, but also provide robust, complete tests of technologies that are extensible and adaptable to other applications. In addition, the program will select applications as the underlying networking technologies begin to enable the necessary infrastructure. Point design studies can determine the maturity of the suite of required network services.

The requirements for selected applications are the following:

1. The application domain is an important Federal mission and is recognized by the public as important.
2. The application demonstration requires high performance internetworking technologies and services that will result from NGI R&D.
3. The networking concepts and technologies embodied in the application testbed are extensible to other application domains and scalable to the future commercial Internet.
4. The application community will supply resources for the application specific technologies component of the testbed.

Applications meeting these requirements will be prioritized on the following basis

1. Resource use and benefit
2. The timely availability of advanced network speed or services
3. The degree to which the application stresses the network technologies being developed in other parts of the initiative
4. The degree to which the applicationís technology plan satisfies the NGI technology requirements including the defined approaches for demonstration, evaluation, validation, scalability, and deployment or commercialization of the resulting capability

4.2.5 Applications Support Function

Each candidate application proposal will come with a defined set of support requirements needed from NGI Goals 1 and 2, that is, the site connectivity and performance requirements needed, as well as the qualities of service, other network services, security services, and middleware products required for the application to be implemented. This candidate set of requirements will be iterated with the tentative plans being developed by Goals 1 and 2, utilizing an Applications Support Function defined by the NGI Implementation Team.

This Applications Support Function will consist of technical support staff, Web site information resources, agency resources, and defined procedures by which a preliminary NGI Application Support Plan and Agreement (pNASPA) is developed. After this pNASPA is iterated through the involved agencies, affinity groups, and NGI support functions, the process results in a final application-specification package that is formally approved or disapproved by the NGI program. At the time of approval, the “p” is removed and the NASPA becomes the support plan and agreement defining each NGI agency and external organization role in supporting the application through the implementation, integration, demonstration, and technology transfer phases.

4.2.6 Applications Affinity Groups

Affinity groups are established where a number of interests need to be coordinated. Applications which are important to only one agency such as national security for DoD and aerospace engineering for NASA, are handled within that agency and do not require affinity group coordination.

Affinity groups have one of two characteristics. First are the *disciplinary affinity* groups (see fig. 8). These are a collection of end-user organizations which share common interests such as health care, education, or environment. They collaborate because they recognize that their applications have a great deal in common; and that by collaboration each will realize its goals more efficiently and effectively.

Second, are the *technology affinity* groups. These groups have as their mission to coordinate and develop the middleware or tools which lie among the network and the applications. For example, many applications require the ability to collaborate over the NGI. Therefore, a collaborative tools affinity group has been established to minimize duplication and to maximize efficiency. They are to ensure that collaboration tools developed by one application are useful to all.

Also, there are two technology groups whose purpose is to coordinate with Goals 1 and 2. Their members are drawn from the other groups as needed.

Disciplinary affinity groups	**Applications technology affinity groups**
Health care	Collaborative technologies
Environment	Distributed computing
Education	Digital libraries
Manufacturing	Remote operations
Crisis management	Security and privacy

Basic science

Federal information services

Each area is reviewed by an expert working group called an affinity group to develop a cross discipline/technology matrix as shown below.

	Applications Technology Affinity Groups				
Discipline affinity groups	Collaboration	Digital libraries	Distributed computing	Privacy/ security	Remote operations & simulations
Basic science	X	X	X	X	X
Crisis management	X	X	X	X	X
Education	X	X	X		X
Environment	X	X	X		X
Federal information services	X	X	X	X	
Health care	X	X	X	X	X
Manufacturing	X	X	X	X	X

Figure 8. Sample NGI cross discipline affinity groups and corresponding applications technology affinity groups Matrix.

The chairs of the affinity groups will work together to provide unified recommendations and requirements to Goals 1 and 2. Key chairs and experts will be part of the implementation team established in the Coordination section of the Executive Summary. Unified recommendations of the affinity groups will be considered authoritative as to the degree to which a candidate application fulfills the NGI criteria and so will be given considerable weight in site selection and service phasing decisions.

Discipline affinity groups have a number of common activities and outputs including the following:

1. Identify applications within the scope of the group which require NGI bandwidth or services. Encourage agencies to think “out of the box” to develop applications which, while impossible without the NGI, will improve mission success using NGI technologies. Endorse those applications. Develop solicitation language to be used by agencies to reach out to the broad community to accomplish key identified applications.
2. Characterize candidate applications according to the prioritization schemata (NGI criteria, resource requirements, and technology expectations).
3. Identify needs common to other affinity groups (collaborative tools, distributed computing, etc.) and identify individuals representing the applications to bridge among the discipline and technology groups.
4. Coordinate and transfer common elements among applications. These common elements may be unique to this discipline or may be shared with other affinity groups.
5. In coordination with other affinity groups, identify NGI capabilities required by applications to produce a superset of the most stringent requirements for the

applications. Interaction with Goals 1 and 2 is intended to be part of the requirements setting process.

6. Harmonize the implementation schedules of the applications and of Goals 1 and 2.
7. Work with Goals 1 and 2, and with the agencies to secure network access, bandwidth, and services for important leading edge NGI applications. Support the development of Multiagency agreements and NGI Application Support Plan and Agreements (NASPA).

Below are the statements of the purpose of each of the affinity groups. The above activities are common to all discipline groups; the only activities listed below are those that are unique to specific groups.

4.2.6.1 Disciplinary Affinity Groups

Health Care

Agencies: NIH, AHCPR, NASA

Scope: The scope of the health care affinity group includes NGI applications which have relevance to the areas of clinical medicine, patient health status, public health, and the health education of professionals as well as the lay public. Many of these NGI applications fall into the categories of advanced telemedicine, telehealth, and distance learning or control applications. They would generally require the transfer of many gigabits of data in close to real time such as magnetic resonance imaging (MRI) or positron emission tomography (PET) scan studies.

Other applications require the transfer of smaller amounts of data but with QoS considerations such as very tight control of latency or jitter (e.g., echocardiography, angiography, nystagmus gait analysis and functional MRI). Still other applications require the transfer of very detailed images within a reasonable time such as pathology and mammography.

Many of these applications require the retrieval of reference multimedia data from libraries. The availability of the NGI will lead to a whole new set of telepresence applications which are based on the ability to control, feel, and manipulate devices at a distance. Applications already being developed include remote microscopy for pathology, remote monitoring, and control of devices for home health care. Eventually, these advances may even lead to telesurgery. All health care applications have a strong security and confidentiality component.

Activities: The health care affinity group will focus on outreach to the health care community, much of which is just discovering the advantages and efficiencies afforded through the use of advanced communications technologies, such as the Internet. The largest use of the Web is the search for health information. Getting medical practices connected, especially in rural areas, is still a formidable problem. The health care affinity group will advocate progress in these areas while encouraging, publicizing, and showcasing advanced patient care applications which make use of the evolving NGI capabilities.

Potential Applications: Telemedicine applications involving highly detailed images (pathology and mammography), real time interactivity (MRI or PET scan studies) and multimedia reference libraries including patient records. Telepresence applications for remote manipulation, feel, and control, as well as home health care.

Environment

Agencies: DoE, NASA, NSF, NOAA, EPA, DoI, DoAg

Scope: Environmental science and services are advancing rapidly based on two related facts. First, our ability to observe the environment is expanding rapidly. Earth based radars, advanced satellite observing systems, and ocean tomography each provide dramatic increases in available environmental data approaching a petabyte per year in the aggregate. Second, advancing computer technology is leading to larger models with large data outputs which can often be best analyzed by advanced immersive environments. These both require the speed and services of the NGI to be successful. Environmental applications span the ocean, atmosphere, and land from short term weather forecasts to long term climate predictions. This includes environmental modeling for hazardous spill/release tracking, pollution transport, and management of Federal environmental resources

Activities: In addition to the common activities identified above, the environment affinity group will focus on those aspects of the NGI which will improve data and information sharing, and model development across the many agencies working in this domain. The group will coordinate closely with the crisis response affinity group to ensure that environmental observations and models are useful in crisis situations. It will also coordinate closely with the basic sciences group to ensure that collaboration and remote operation technologies are shared across both domains.

Potential Applications:
Climate collaboratory; Chesapeake Bay virtual environment; distributed modeling laboratory for mesoscale meteorological studies; and real time environmental data via the NGI.

Education

Agencies: DoEd, NASA, NSF

Scope: Numerous educational initiatives are underway in the Department of Education (DoEd) including the Technology Literacy Challenge Fund, Technology Innovation Challenge Grants, and several activities which target universal access in education. These activities, coupled with the Federal Communications Commission (FCC) discounts for schools and libraries which total $2.25 billion annually lay the foundation for immediate connectivity to the classroom and for the development of innovative education content for the next generation of networked education applications. NGI, working through the agencies and the CCIC R&D Subcommitteeís Education, Training and Human Resources Working Group (ETHR), will build on this foundation to reach out and develop exciting new applications.

Activities: Most education is carried out at the state and local level, hence outreach to the broadest possible community will be a priority for this group.

Potential Applications:
Distance learning; universal access.

Manufacturing

Agencies: NIST, DoE, NASA

U.S. manufacturers are implementing new organizational models, new engineering and manufacturing processes; incorporating new materials; and adopting new quality methods in efforts to achieve best-in-class performance. Implementing any of these efforts in a single manufacturing facility affects the information technology infrastructure of that facility. When these efforts are implemented over a geographically distributed enterprise, supply chain, or virtual enterprise, the global information network becomes the constraining factor determining the degree to which information-intensive efforts are realized. Fully realizing these efforts depends on the successful deployment of a network infrastructure which provides reliable data transfer, deterministic propagation delay, privacy, and a variety of data capacities.

Scope: The manufacturing affinity group will focus on applications representative of business activities which manufacturers will expect to be enabled by the NGI. Typical activities could include establishment of virtual organizations, collaborative product/process design and engineering, remote equipment control and monitoring, managing distributed workflows, accessing distributed manufacturing data, and sharing software environments.

Activities: The principal activities of the manufacturing affinity group will be to identify agency manufacturing applications, analyze manufacturing applications for network and services requirements, and collate those requirements for consideration by other NGI working groups. In addition, it will help to identify commonalities among applications to enable leveraging among them.

Potential Applications:
Characterization, remote access, and simulation of hexapod machines; diesel combustion collaboratory; telerobotic operation of scanning probe microscopes.

Crisis Management

<u>Agencies</u>: FEMA, NOAA, NASA, DoD/DARPA, USGS, DoE, DoI, CEOS, G-7/GEMINI, GDIN NOAA/NESDIS

*<u>Scope</u>***:** Numerous crisis (disaster) programs and projects are underway within the U.S. In many Federal agencies there are projects related to some aspect of crisis management or disasters, including prediction, forecasting, monitoring, response, assessment, mitigation, relief coordination, and intervention/ assistance. Projects are being developed and funded to meet an often specific (narrow) objective and use limited resources (i.e., those of the particular agency or in some cases several cooperating agencies). This is in part because today's telecommunications speeds and services cannot meet the broader requirements that many of these projects could implement.

Regardless of current limitations, many of these applications are planning for NGI-type services. They have goals to expand the accessibility and increase the utilization of the vast array of technology, services, and information currently within the Federal government (as well as those technologies and services being planned) to significantly improve the nation's forecasting, preparedness, and response to crisis management and disaster situations.

<u>Activities</u>: The crisis management affinity group has perhaps the most difficult networking requirements. This will require tight collaboration with the engineering teams. The wide variety of governmental organizations involved in crisis management requires close coordination with the involved organizations, and with all national information assets.

<u>Potential Applications</u>:
Collaboration under crisis conditions; data access and fusion under crisis conditions; security and privacy policies and enforcement ; nomadic computing and network management in a crisis.

Basic Science

<u>Agencies</u>: DoE, NSF, NOAA, NASA, NCRR

<u>Scope</u>: The discipline of conducting scientific research has been undergoing a sometimes subtle, but accelerating, evolution. Advances in computing and communications technology are being assimilated into the scientific environment and are bringing changes in how scientists interact with their peers, their data, and their facilities. An era in which new paradigms of science are possible is approaching. It is important to enable applications that clearly demonstrate that it is possible to do science in new ways, through NGI technology, ways that would not have been possible otherwise.

<u>Activities</u>: The Basic Science Group will focus on applications that provide insights into fundamental science and associated phenomena. Criteria for selection of applications specific to this discipline will be developed.

Organization of the basic science affinity group is self-selecting in the initial formation, and membership is derived from representatives of the agency mission critical and signature applications.

Potential Applications:
Diesel combustion collaboratory; materials microcharacterization collaboratory; advanced computing initiative; the Oak Ridge/Sandia distributed computing project.

Federal Information Services

Agencies: Various

Scope: These applications include the full range of advanced information services of the government. They potentially span the full range of governmental levels. This group will link the providers of Federal information services and the networking research community to provide the early adopters of the information services community with the opportunity to help guide the services of the NGI and to explore the new services these advanced capabilities enable. A few applications require the highest bandwidths, but many in this domain require the advanced services which support security, privacy, collaboration, and distributed knowledge discovery.

Activities: Most Federal information service providers are fully occupied implementing the technologies available today, hence only a few are prepared to participate in the NGI. This affinity group will need to strongly market the program to the intended community.

Potential Applications: Storage, archival, and information access; information integration; data mining; electronic commerce.

4.2.6.2 Technology Affinity Groups

Technology affinity groups have a number of common activities and outputs, including the following:

1. Identifying needs common to other affinity groups (collaborative tools, distributed computing, etc.) and identifying individuals representing the applications to bridge among the discipline and technology groups.
2. Coordinating and transferring common elements among applications. These common elements may be unique to this technology or may be shared with other affinity groups.
3. In coordination with other affinity groups, identifying NGI capabilities required by applications to produce a superset of the most stringent requirements for the applications. Interaction with Goals 1 and 2 is intended to be part of the requirements setting process.

4. Harmonizing the implementation schedule of the technologies, the applications, and the Goal 1 and 2 implementation schedules.

Collaborative Technologies

Agencies: NIST, DoE, NASA

Scope: The use of collaborative technologies is critical to the success of a broad range of government projects. These technologies include, for example, network based videoconferencing, shared documents and notebooks, shared databases, and remote access to shared computers and research facilities. Together, these technologies permit scientists, engineers, and administrative staff to work together on projects without regard to physical location.

The collaboration working group will provide guidance on the technology needs for collaborative applications and on the availability of existing and future tools to meet those needs. The group will be composed of people who are involved in collaborative applications in many Federal agencies.

Activities: In particular, the working group will be concerned with the following:

1. Evaluate the needs for collaborative technologies across a broad range of agency applications, and prepare a summary of requirements that are likely to be met by common solutions and those that are unique to specific applications.
2. Evaluate the implications of the collaborative technologies across the applications for networking infrastructure, management, and services.
3. Summarize the current state of collaborative technologies available commercially and from the private sector, and assess the ability of those technologies to meet the common needs of agency applications.
4. Determine where current and planned R&D will address needs that are not satisfied by existing collaboratory tools, and make recommendations for appropriate R&D that is likely to provide solutions that meet agency needs.

In carrying out this review, the working group will maintain close communication with all applications groups, as well as with appropriate technology working groups such as security, distributed computing, and remote access. The group will also work closely with groups working on NGI Goals 1 and 2 to ensure that infrastructure and research plans take into account the needs of collaborative applications.

Distributed Computing

Agencies: DoE, NSF, NASA

Scope: The Internet is more than just a communication medium: it is also a means by which computation can be integrated with human activities. Many proposed NGI applications depend on the ability to access, in a coordinated fashion, remote computers, storage systems, databases, scientific instruments, advanced display devices, and other resources. Distributed computing technologies will allow this access to occur in a straightforward, efficient, and secure way.

Distributed computing is about more than remote procedure calls (RPCs). The scale and heterogeneity of the Internet means that applications also need to be able to locate and schedule resources (including networks); determine properties of resources; configure resources and computations; support diverse communication mechanisms, including RPC, message passing, streaming video, and multicast; access and manipulate distributed data stored in diverse forms; monitor and manage computations, etc. Emerging computational models based on mobile agents introduce additional issues.

A number of approaches have been proposed that seek to provide these capabilities. However, each seems a partial solution, and none has been tested on the scale of the national scale “computational grids” that we expect to evolve in the future. Perhaps as a result, we see little agreement in specifics, and significant obstacles to the large scale experiments that might accelerate progress.

Activities:

1. The working group will identify and characterize the many different classes of expected NGI applications and their distributed computing requirements. These requirements will encompass a variety of issues including resource location, resource scheduling, process management, configuration, communication, data access, security, QoS management, mobile code.

2. (Proposed) Further, the working group will evaluate current distributed computing approaches in terms of their ability to provide the services required by the extremely diverse, widely distributed, and potentially autonomous applications expected to provide and manage the content of the NGI.

3. (Possible) The working group will also provide recommendations for courses of action that will speed the development and deployment of the distributed computing services identified in Activity 2, that are needed by widely distributed and diverse applications identified in Activity 1.

Digital Libraries

Agencies: NSF, NASA, NIH

Scope: This topic has evolved rapidly to become a rich field yet one that has not fully matured. Even the term “digital library” is being defined and redefined as capabilities not available in traditional libraries are being developed and their effect better understood. For example, the contour map on the page of a book can be turned into a “fly-over,” and information from a variety of sources about one place—be it a place on the map or in the brain—can be integrated into a three-dimensional picture from which future (or past) behavior can be simulated.

As digital libraries technologies mature and are disseminated, and as the size of the national, in fact the worldwide, digital library grows, every citizen will become able to take out a “book” from any library or interactively visit any museum. They are already being used in crisis response; these libraries contain images of how things were prior to natural disasters that can be quickly accessed by response teams. Their application will span and expand human knowledge. These uses will require NGI speed and extent.

Digital libraries require distributed mass storage systems for storing their repositories; high performance networking for users to access content, which often takes the form of multiple multimedia objects; and high performance computing to manipulate the data (for example, to move around three-dimensional data sets such as environmental or biomedical models).

Activities: Future activities will build on the ones that have been conducted over the last several years. Indeed, one major reason for the rapid developments in digital libraries is the NSF/DARPA/NASA Joint Digital Libraries initiative. Six university led consortia that include libraries, museums, publishers, schools, and computing and communications industry companies are conducting R&D in digital libraries technologies in this 1995 to 1998 effort. The NGI digital libraries affinity group will coordinate these activities with the needs of the NGI applications to ensure that digital library technologies will operate over the NGI and integrate with the applications.

Remote Operations

Agencies: NIST, DoE, NASA, NIH, NSF

Scope: Many proposed NGI applications involve the remote monitoring of a machine, process, environment or crisis situation. The term “remote operations” implies the ability of the remote user to effect a change in what is being monitored and see results in a timely fashion. This ability to effect change can range from the simple setting of a few key parameters to taking remote control of some device at a very low level such as force feedback. The technologies and abilities being developed for NGI come into play at all these levels.

Privacy and security issues are key when operating remote, unique devices or gathering proprietary information on the performance of a product. The development of collaboratories often includes remote operation of devices at various levels of control. The high bandwidth and deterministic nature of communications using NGI network technology is required for remote control of high-control bandwidth applications such as force reflection (remote surgery, bomb dismantling) or remote diagnostics or calibrations.

An initial taxonomy of applications based on latency requirements is shown below. The Level 3 and 4 applications are achievable using current Internet technology, but suffer from bandwidth limitations and the nondeterministic latencies. Level 1 and 2 applications are expected to require deterministic communications. It is envisioned that these applications will first request and then determine the end-to-end network latency, incorporating this delay into their control law calculations. The maximum delay tolerable is a function of the control laws being used and the physical plant being controlled. Certain Level 1 applications may require a minimum number of switch delays or be otherwise limited by speed of light considerations to local users rather than coast-to-coast operation.

Level	Typical application	Latency sec
4	Weather station parameter settings	> 100
3	Chemical plant process control	10 - 100
2	Machine tool control	1 - 10
1	Force reflection for remote surgery	0.1 - 1

Activities: The remote operations affinity group will initially focus on exploring the minimum requirements of testbed applications and will seek to encompass all levels of potential remote operations. In conjunction with this effort and in concert with the other affinity groups, we will seek to develop a taxonomy of remote operations with respect to their NGI services and mechanisms for the application to request and verify that the proper services (bandwidth, latency, user validation) are in place before allowing the remote use of what are often unique national resources.

Security and Privacy

Agencies: DoE, NIST, NASA

Scope: Application level security plays a much broader role than does infrastructure security, IP-level security, and system level security. Although the lower levels of security are essential to protecting the computing and communications infrastructure, Application level security is as much (or more) about enforcing agreements among “legitimate” users as it is about protection.

Application level security is the mechanism that will enable widely distributed enterprise. It must provide, for example, for expressing use conditions on data, services, and resources; expressions of authorization and attributes; and payment mechanisms. It also must provide the mechanisms for fine grained protection of these assets, assuming that the infrastructure is secured.

Several mechanisms—security architectures, infrastructure, and technology—have been proposed for Application level security, with naming authority and third party trust mechanisms, public key cryptography, and cryptographically signed certificates as general, common threads, but with limited, or only emerging, agreement on specifics.

Activities: The following activities will be pursued in the context of and in coordination with the broad community outside NGI this is concerned with privacy and security, including the CCIC R&D Subcommitteeís High Confidence Systems Working Group.

1. The Working Group will identify and characterize the many different classes of expected NGI applications and their security requirements. This will encompass a variety of issues including, for example, limited authority delegation in global financial systems; protecting personal privacy in the face of comprehensive, online personal attribute databases; legally mandated access controls for medical records; securing and enforcing resource use agreements in globally distributed meta-computers.

2. (Proposed) Further, the Working Group will evaluate current security approaches in terms of their ability to provide the services required by the extremely diverse, widely distributed, and potentially autonomous applications expected to provide and manage the content of the NGI.

3. (Possible) The Working Group will also provide recommendations for courses of action that will speed the development and deployment of the security services identified in Activity 2, that are needed by widely distributed and diverse applications identified in Activity 1.

4.3 CANDIDATE APPLICATIONS

Suggested demonstrations cover a wide range of capabilities: from time critical applications such as crisis and national security responses to broad collaboration in areas as diverse as health care, education, and research. Telemedicine extends collaboration adding robustness, security, and reliability.

Application testbeds serve as platforms for proof-of-concept demonstrations. They tie together networking technologies, test the completeness of NGI protocols, and force the technologies to operate in real situations. By forcing technologies to work together in complex situations, applications stress the cooperability and interoperability of the developing suite of advanced networking services. In addition, effective demonstrations showcase new network capabilities, resulting in new acceptance and even enthusiasm for these important advances.

4.3.1 Potential Applications

Health Care: Doctors at university medical centers will use large archives of radiology images to identify the patterns and features associated with a particular disease. With remote access to supercomputers, they will also be able to improve the accuracy of mammography by detecting subtle changes in three-dimensional images.

Crisis Management, Crisis Response: Crisis managers will access a wide range of information under the most difficult and unpredictable circumstances. Networks will be self-configuring to enable rapid return of services after a disaster; and information from multiple levels of government and the public sector will be immediately available. The results of remote models will be available for failure diagnosis and prediction of effects during natural or man-made disasters.

Education, Distance Education: Universities are now experimenting with technologies such as two-way video to remote sites, VCR-like replay of past classes, modeling and simulation, collaborative environments, and online access to interactive, multimedia instructional software. Distance education will improve the ability of universities to serve working Americans who want new skills but who cannot attend a class at a fixed time during the week.

Basic Science, Scientific Research: Scientists and engineers across the country will be able to work with each other and access remote scientific facilities as if they were in the same building. “Collaboratories” that combine videoconferencing, shared virtual work spaces, networked scientific facilities, and databases will increase the efficiency and effectiveness of our national research enterprise.

Environment: Climate Research: Scientists, researchers, and policy makers will be able to examine the effects of proposed actions on the long term evolution of our environment. Models will become available to and usable by all interested users.

Health Care: Biomedical Research: Researchers will be able to solve problems in large scale DNA sequencing and gene identification that were previously impossible, opening the door to breakthroughs in curing human genetic diseases.

Environment: Environmental Monitoring: Researchers are constructing virtual worlds to model and monitor defined ecosystems. For example, one project models the Chesapeake Bay ecosystem, which serves as a nursery area for many commercially important species.

Manufacturing: Collaborative engineering, distributed data sharing, and teleoperation of unique manufacturing resources will dramatically reduce the time required to develop new, higher quality products in distributed enterprises and in virtual enterprises. Enhancing manufacturing applications such as design, analysis, modeling, simulation, virtual reality, equipment/process control, and monitoring with NGI-based capabilities will enable system-wide improvements for geographically distributed manufacturers, suppliers, and customers.

4.3.2 Initial Candidate Applications

The following initial candidate applications have been developed within individual agencies. As the program progresses and the number of applications increases five-fold or more, applications which cut across multiple agencies and even multiple sectors will be encouraged.

DoE Applications: Diesel combustion collaboratory; materials microcharacterization collaboratory; accelerated strategic computing initiative; accelerated strategic computing initiative; the Oak Ridge/Sandia Distributed Computing Project; data access and analysis of massive datasets for high energy and nuclear physics.

NASA Applications: NASA instrument quality assurance; instrument support terminal; NASA echocardiography; distributed image spreadsheet ; collaborative simulation; virtual simulation.

NIH Applications: Radiology consultation workstation; distributed positron emission tomography (PET) imaging; Real time telemedicine; high resolution imaging telemedicine; remote control telemedicine; medical image reference libraries.

NIST Applications: Telerobotic operation of scanning tunneling microscopes; characterization remote access and simulation of hexapod machines.

NOAA Applications: Crisis management – a collection of generic requirements advanced numerical weather forecasting.

NSF Applications: Chesapeake virtual environment; distributed modeling laboratory for mesoscale meteorological studies; real time environmental data via the NGI.

4.4 MILESTONES

Fiscal Year	
4Q1997	In conjunction with the CCIC Applications Council, identify high priority mission applications which require advanced networking technologies and services
2Q1998	Attract key application mission partners
4Q1999	Demonstrate applications using the first advanced networking technologies (IPv6, ATM, QoS)
1Q1998, 1999, 2000	Determine technical dependencies from Goals 1 and 2; match applications demonstrations with those deliverables
4Q2000	Demonstrate enhanced applications as networking technologies evolve (security, nomadicity)
4Q2002	Leverage experience gained from early application demonstrations to develop and to demonstrate more complex applications

Metrics
Number of applications which use 100x speed
Number of applications which use 1000x speed
Number of applications which use Goal 1 services:
QoS, nomadicity, adaptive network management, electronic commerce level security, and patient record level privacy.

5. MANAGEMENT PLAN

The NGI Program will be coordinated within the framework of the National Science and Technology Council (NSTC). The Committee on Computing, Information, and Communications (CCIC) will be responsible for the overall high level NGI strategy. The Computing, Information, and Communications (CIC) R&D Subcommittee is responsible for coordination across program component areas. The Large Scale Networking (LSN) Working Group is responsible for the implementation strategy of the NGI. A small, integrated NGI Implementation Team will take primary responsibility for implementing the approved plans under the direction of the LSN Working Group.
In particular, the NGI Implementation Team will:

1. Contain one member from each of the funded agencies plus an applications advocate who will provide linkage to NGI applications partners and to the CCICís Applications Council
2. Use advanced networking and computing for effective coordination and communications
3. Answer to the LSN Working Group as a team (and to agencies as individuals)
4. Operate as an integrated project team for the overall NGI initiative
5. Be jointly responsible for execution of approved implementation plans, initiative management and evaluation, and other activities as required for successful implementation
6. Establish contributing partnerships and relationships
7. Recommend funding mechanisms and serve appropriately in the selection process.

The directly funded agencies will, of course, also participate in the oversight of the NGI implementation team by, among other things, the approval processes required to expend agency resources in support of the NGI initiative.

6. SUMMARY

The Next Generation Internet initiative will build the foundation for powerful and versatile 21st century networks. The NGI partnership among government, industry, academia, and the general public will bring these diverse talents into focus to solve the problems that have arisen from the growing complexity and magnitude of requirements being placed upon the Internet. Without this partnership, solutions will not be forthcoming for at least another decade, and America's technological leadership position will be at risk.

Goal 1 of the NGI initiative is to research, develop, and experiment with advanced network technologies that will provide dependability, diversity in classes of service, security, and real time capability for such applications as wide area distributed computing, teleoperation, and remote control of experimental facilities. These activities focus on network growth engineering, end-to-end QoS, and security.

Accompanying the development of advanced network technologies is Goal 2 of NGI's initiative, development of the next generation network fabric. This effort will overcome today's speed bumps which slow end-to-end usable connectivity; the slowing is caused by incompatibilities in switches, routers, local area networks, and workstations. Two thrusts within this goal are planned: First, construction of a high performance distributed laboratory consisting of the 100 NGI sites at universities, at Federal research institutions, and at other research partners at speeds in excess of 100 times that of today's Internet. This laboratory will be large enough to provide a full system, proof-of-concept testbed for hardware, software, protocols, security, and network management required by the commercial NGI. Second, development of ultrahigh speed switching and transmission technologies and end-to-end network connectivity at more than 1 gigabit per second. Such networks will be high risk, pioneering networks limited to 10 NGI sites at speeds 1,000 times faster than today's Internet.

These two goals—experimental research of advanced network technologies and development of the next generation network fabric—will provide the basis for terabit-per-second networks operated by appropriate network management and control providing guaranteed end-to-end QoS.

Finally, to test the advanced capabilities of the first two goals, Goal 3 will demonstrate a selected number of applications requiring these capabilities over the NGI network(s). Procedures have been established to ensure that selected applications provide robust, realistic, complete tests of technologies that can be extended and adapted to other applications. Initial applications are being chosen from the Federally focused applications in appropriate technology classes, such as digital libraries, remote operation of medicine, and crisis management.

Participating agencies coordinate the NGI joint agency effort, and since most of the funds will be provided by the applications, leadership is provided by the domain-specific affinity groups. Agencies conduct their own calls for research and coordinate with other agencies to review the proposals received. An NGI implementation team coordinates research agendas across all goals.

The multi-agency NGI initiative—a solid partnership with industry, academia and the American public—thus provides the catalyst for the development of high performance, secure, reliable networks of the future and ensures continued U.S. dominance in the world's high technology networking arena.

7. APPENDIXES

7.1 APPENDIX A

ACRONYMS AND ABBREVIATIONS

A	
AADS	Ameritech Advanced Data Services
AAI	ACTS ATM Internetwork
ABR	Available bit rate
ACTS	Advanced Communications Technology Satellite
AHCPR	Agency for Health Care Policy and Research
ANL	Argonne National Laboratory
API	Application Program Interface
ATM	Asynchronous Transfer Mode
B	
BAA	Broad Area Announcement
BGP	Border Gateway Protocol
BIT	Broadband Information Technology
C	
CA*net	Canada’s Research and Engineering Network
CAN	Cooperative Agreement Notice
CBQ	Class Based Queuing
CCIC	Committee on Computing, Information and Communications (White House Office of Science and Technology Policy)
CEOS	Committee on Earth Observation Satellites
CERN	European Laboratory for Particle Physics
D	
DARPA	Defense Advanced Research Projects Agency
DFN	Deutsche Forschungsnetz (German Research Network)
DISC	DREN Intersite Service Contract
DN-VBN	DoE/NGI Virtual Backbone Network
DoAg	Department of Agriculture
DoD	Department of Defense
DoE	Department of Energy
DoEd	Department of Education
DoI	Department of the Interior
DREN	Defense Research and Engineering Network
DS3	Digital Signal 3 (44.7 Mbps)
E	
EOS	Earth Observing Satellite
EOWG	Engineering and Operations Working Group (of the FNC)
EPA	Environmental Protection Agency
EPSCoR	Experimental Program to Stimulate Competitive Research
ESnet	Energy Sciences Network (DoE)
ETHR	Education, Training and Human Resources Working Group
F	
FCC	Federal Communications Commission
FEMA	Federal Emergency Management Agency
FNC	Federal Networking Council
FY	Fiscal Year
G	
Gbps	Gigabits per second (10^9)

GIBN	Global Interoperability Broadband Network
GDIN	Global Disaster Information Network
GigaPOP	Gigabit Points of Presence
H	
HDR	High Data Rate
HFC	Hybrid Fiber Coax
HIPPI	High Performance Parallel Interface
HPCC	High Performance Computing and Communications
HPCMP	High Performance Computing Modernization Program
I	
IC&V	Intelligent Collaboration and Visualization
IETF	Internet Engineering Task Force
I-NNI	Network to Network Interface
IP	Internet Protocol
IPSec	Internet Protocol (Secure)
IPv4/ IPv6	Internet Protocol, versions 4 and 6
ISAKMP	Internet Security Association Key Management Protocol
ISP	Internet Service Provider
IT	Information Technology
J	
JET	Joint Federal Engineering Team
JLAB	Jefferson Laboratory
K	
kbps	kilobits per second (10^3)
L	
LAN	local area network
LSN	Large Scale Networking
M	
Mbps	Megabits per second (million)
MRI	Magnetic Resonance Imaging
msec	millisecond
MTPE	Mission to Planet Earth
N	
NAP	Network Access Point
NASA	National Aeronautics and Space Administration
NASPA	NGI Application Support Plan and Agreement
NCRR	National Center for Research Resources
NCSA	National Center for Supercomputing Applications
NESDIS	National Environmental Satellite, Data and Information Service
NGI	Next Generation Internet
NGI IT	Next Generation Internet Implementation Team
NGIX	NGI eXchange
NIH	National Institute of Health
NIST	National Institute of Standards and Technology
NLANR	National Laboratory for Applied Network Research
NLM	National Library of Medicine
NNI	Network to Network Interface
NOAA	National Oceanic and Atmospheric Administration
NOC	Network Operations Center
NPD	Network Probe Daemon
NREN	NASA Research and Education Network
NSA	National Security Agency
NSF	National Science Foundation
NSTC	National Science and Technology Council

O

OC-3	Optical Carrier-3 (155 megabits per second)
OC-12	Optical Carrier-12 (622 megabits per second)
OC-48	Optical Carrier-48 (2.5 gigabits per second)
P	
PACI	Partnerships for Advanced Computational Infrastructure
PET	Positron Emission Tomography
PET	Programming Environment and Training
PKI	Public Key Infrastructure
pNASPA	preliminary NGI Application Support Plan and Agreement
P-NNI	Peer Network to Network Interface
PoA	Point of Attachment
POP	Points of Presence
PPP	Point to Point Protocol
PVC	Permanent Virtual Circuit
Q	
QoS	Quality of Service
R	
R&D	Research and Development
R&E	Research and Engineering
RED	Random Early Drop
RF	Radio Frequency
RPC	Remote Procedure Call
RSVP	Resource Reservation Protocol
RTP	Real Time Protocol
S	
SBIR	Small Business Innovative Research
SDSC	San Diego Supercomputing Center
SONET	Synchronous Optical NETwork
STAR-TAP	Science, Technology and Research Transit Access Point
STTR	Small Business Technology Transfer
SuperNet	Terabit Research Network (DARPA)
T	
Tbps	Terabits per second (10^{12})
TCP	Transport Control Protocol
U	
UCAR	University Corporation for Atmospheric Research
UNI	User Network Interface
URL	Uniform Resource Locator
USGS	U.S. Geological Survey
V	
VBN	Virtual Backbone Network
vBNS	Very High Speed Backbone Network Service (NSF)
W	
WAN	Wide Area Network
WDM	Wavelength Division Multiplexing

7.2 APPENDIX B

UNIFORM RESOURCE LOCATORS

Agency or Institution	Acronym	URL
Dept. of Defense/ Defense Advanced Research Projects Agency	DoD/ DARPA	http://www.uncle-sam.com/defense.html http://www.ito.darpa.mil/ResearchAreas.html
Dept. of Energy	DoE	http://www.doe.gov/
DoE's Energy Science Network	ESnet	http://www.es.net
Federal Networking Council	FNC	http://www.www.fnc.gov/
Internet 2	I2	http://www.internet2.edu/
NASA Research and Education Network	NREN	http://www.nren.nasa.gov
National Aeronautics and Space Administration	NASA	http://www.nasa.gov/
National Institute of Health	NIH	http://www.nih.gov/
National Institute of Standards and Technology	NIST	http://www.nist.gov/
National Oceanic and Atmospheric Administration	NOAA	http://www.noaa.gov/
National Science Foundation	Connections Program vBNS Network	Http://www.cise.nsf.gov/ncri/hp-connections.html http://www.vbns.net
National Science Foundation	NSF	http://www.nsf.gov/
Next Generation Internet	**NGI**	**http://www.ngi.gov**

7.3 APPENDIX C

DISTINGUISHING CHARACTERISTICS OF REVOLUTIONARY APPLICATIONS

The following is a description of the functional NGI application requirements. These requirements are to be met by an appropriate combination of the applications themselves, technology affinity groups, and the capabilities of Goals 1 and 2.

1. **Security**. Telemedicine and electronic commerce, for example, will rely on the capability to maintain privacy and the confidentiality and integrity of personal data.
2. **Data Sharing**. Digital libraries, other science and technology information banks, etc., will be required for network based applications such as federated genome data bases, crisis response, and Earth Observing Satellite (EOS) data used throughout the Space and Earth Sciences community.
3. **Software Sharing**. Scientists at different locations will need the capability to conveniently share software that supports data analysis, visualization, and modeling to all manner of remote collaborations.
4. **Controlling Remote Instruments**. Communicating with distant fellow workers is required for aerodynamic design and for using such remote science facilities as an advanced light or photon source across a network.
5. **Visualization**. Remote visualization technology is important for seeing what is being controlled at a remote facility or for viewing the results of computational simulations. Advanced visualization technologies such as network integrated, immersive virtual reality devices will be needed to allow multiple design or experimental teams to work together across distances to simultaneously observe or analyze data, images, etc.
6. **Scalability**: Network technologies used by wide area applications must be able to be scaled up to support applications at the national level far better than is possible today.
7. **High end Computation and Computing Resources**. Testbeds will need to integrate supercomputers and computational technologies for a number of reasons. In remote experimentation, supercomputers may be used for real time diagnostics to ensure that devices are performing within specification. In other forms of telescience, supercomputers may be used for instrument recalibration or for real time modeling of experimental data.
8. **Self-Organizing Networks**. This capability provides self-adaptation when the physical configuration or requirements for network resources have changed. Crisis management requires the ability to establish or reestablish networks in the field among managers, action agents (such as police, fire, health care), and situation-specific information.
9. **Nomadicity**. The ability to move resources as needed will become increasingly important. It will include “mobility of access rights” so the network will know

how to treat a new resource. This may range from full rights to complete denial of access.

10. **Rapid Resource Discovery Capability**. Currently, network administrators painstakingly document resources, assign rights, and monitor use. In the future, everyone will require the ability to discover network resources as needed. The most extreme case will be during the response to a natural disaster or other crisis.

11. **Portability and Interoperability of Applications**. As networking and computing become more ubiquitous, work will increasingly be accomplished only with the end-user application requiring the idiosyncrasies of networks and computers to be transparent to users.

12. **Virtual Subnetworking**. This provides the ability to establish specialized communities of interest which may be a group of researchers collaborating on a climate model, a contractor and subcontractors working on a new product, or a task force developing a new policy.

13. **Ease of Use**. At heart of future networks will be ease of use. It will be as easy to add resource to networks as it is to plug in a phone today.

14. **Reliability**. When advance networking services are implemented, they will be fragile and suitable only for research, yet the designs must eventually be scalable to full commercial and even military robustness.

7.4 APPENDIX D

COMMUNITY OUTREACH—PROPOSAL DEADLINES

Agency	Type	Announcement date	Proposal due date
DoE	General	October 1, annually	Throughout the year
	Focused	As needed	As needed
NSF	High performance connections	Continuing	July 31 December 31
	Network Development		Throughout the year
	Meritorious Applications	Per appropriate directorate	Per appropriate Directorate
NASA	Technology partnership agreements	As needed	As needed
	Application solicitations	Continuing	January 31 July 31

In addition to the proposal processes above, all agencies will coordinate with their corresponding commercialization and technology development offices to leverage the Small Business Innovative Research (SBIR) and Small Business Technology Transfer (STTR) opportunities for the benefit of the Next Generation Internet goals.